Università Ca'Foscari Venezia

Corso di Dottorato di ricerca

in Economia

Ciclo XXXII

Tesi di Ricerca

Micro Effects of Macro Policies

SSD: SECS-P/06

Coordinatore del Dottorato

ch. prof. Giacomo Pasini

Supervisore

ch. prof. Giacomo Pasini

Secondo Supervisore

ch. prof. Maarten Lindeboom

Dottorando

Francesco G. Caloia

Matricola 956270

Acknowledgment

Being here to write these few lines makes me realize two things. The first is that this fantastic journey, called PhD, has come to an end. Four years spent in three different institutions, that made me know many fantastic people. I grew personally and professionally and, most importantly, I had a lot of fun! The second thing I realize is that I have a lot of people to thank. So, here I am.

I warmly thank Giacomo Pasini, my first supervisor, for starting to think about an ambitious journey for me since day one. I really liked his pragmatic way of supervising me, and the fact that he has continuously pushed me to raise my motivation and goals as time went by.

I thank my second supervisor, Maarten Lindeboom, for opening to me the doors of the VU and for giving me the opportunity to obtain a double degree. He has always allowed me to work free of any concern, and gave me the flexibility and the freedom of pursuing my own interests for the sake of my thesis.

I am extremely grateful to Mauro Mastrogiacomo for the continuous and endless discussions on our common research interests, for the personal support, as well as for all the things he thought me about this country. We really have fun in working together, especially in those close-to-deadline coding sessions on the loads of data we work on. In such a short time he went from being a co-supervisor, to being a colleague, a coauthor, and finally a friend. I'll never be able to give him back not even half of what he gave to me.

I warmly thank Jante Parlevliet, my manager at the Economic Policy department at DNB. She is simply the perfect manager, the one that senses your skills, knows how to enhance them, and knows how to make them useful for the division. Now that she left our department, I am extremely happy to continue our relationship as coauthors.

I am infinitely grateful to my master supervisor, prof. Silvia Muzzioli. She literally used any mean (including phone calls after a job interview I did at a bank) to convince me that the plans I had at the time were wrong for my skills and interests. Without her, I would have never started this fantastic journey that I enjoyed from the first to the last day.

I thank all my colleagues of the Research and the Financial Stability departments of DNB. Thanks to them, I could seriously discuss the policy implications of my work. This has deeply impacted the outcome of this PhD thesis.

I also thank all the special people I met during this journey. The most special thanks goes to Irene, who has been to my side during the entire period of this PhD, both in Italy and the Netherlands. I thank all my other PhD colleagues from Ca' Foscari University of Venice: we have gone through the toughest and the most joyful moments together. The experiences we shared, as well as the drinks and parties, are simply unforgettable. I thank all the countless interns who went through DNB over the past two years for the great fun we had on our breaks and thursday nights out. I also thank all my colleagues of the VU who I always spent a good time with, despite my sudden and irregular appearances.

I thank my mum and dad, to whom this thesis is dedicated. I thank them for their unconditional support and trust in me and in what I wanted for my future. I thank my sister Giulia, I miss her even though I'm sure she is happy that we don't live together anymore, as I can't bother her on a daily basis. I thank my grandparents for their unconditional love, this is for them too. I finally thank all my friends. Despite the distance, things have not changed and they never will with you guys.
All the people I mentioned had a different, but significant role ($p < .01$ ***) in making my last four years the most interesting and most rewarding days I've had in my life.

Contents

Preface . . . v

1 Being in Good Hands: Deposit Insurance and Peers' Financial Sophistication 1

1.1 Introduction . . . 1

1.2 Institutional framework . . . 5

1.3 Data and Descriptive Analysis . . . 6

1.4 Empirical Analysis . . . 9

1.4.1 Wealth allocation under the DGS rules . . . 9

1.4.2 Identification strategy . . . 11

1.4.3 The effect of the DGS reform . . . 13

1.4.4 Consumption cuts . . . 15

1.4.5 Portfolio reallocations . . . 16

1.4.6 Implicit public guarantees . . . 17

1.4.7 Peer effects . . . 19

1.5 Internal validity and robustness . . . 22

1.6 Concluding Remarks . . . 25

2 Leverage Constraints, House Prices and Household Debt: Evidence from the Netherlands 31

2.1 Introduction . . . 31

2.2 Institutional framework . . . 35

2.2.1 Macro-Prudential Policy . . . 35

2.2.2 Mortgage Market . . . 38

2.2.3 Real Estate Market . . . 39

2.3 Theoretical framework 40
2.3.1 Aggregation 44
2.4 Data and descriptive statistics 47
2.5 Empirical analysis 50
2.5.1 The effect of house prices, LTI and LTV limits 50
2.5.2 Bunching at the LTI limit 56
2.5.3 Bunching at the LTV limit 63
2.6 Concluding remarks 69
2.6.1 Conclusions 69
2.6.2 Policy implications 70

3 Staggered Wages, Unanticipated Shocks and Firms' Adjustments 73

3.1 Introduction 73
3.2 Institutional Framework 77
3.3 Data and Descriptive Statistics 79
3.3.1 Collective Agreements Data 79
3.3.2 Matched Employers Employees Data 82
3.4 Empirical Analysis 84
3.4.1 Identification Strategy 84
3.4.2 Results 88
3.4.3 Heterogeneity Analysis 91
3.5 Robustness 93
3.5.1 Agreement vs Start dates 93
3.5.2 Worker-level evidence 96
3.5.3 Parallel Trend test 98
3.5.4 Other robustness checks 99
3.6 Concluding remarks 99

4 The Housing Wealth Effect: a comparative analysis of Italy and The Netherlands 101

4.1 Introduction 101
4.2 Theoretical Framework 104

4.2.1 Exogenous maintenance expenses 105
4.2.2 Endogenous home improvements 106
4.2.3 Quality-specific pricing and renovation costs 109
4.3 Data and Methodology 110
4.4 Descriptive Analaysis 113
4.5 Empirical Analysis 117
4.5.1 Heterogeneity Analysis 123
4.6 Conclusions 124

5 Bibliography **129**

List of Tables

1.1 Descriptive statistics . . . 7
1.2 The general effect of the DGS reform . . . 14
1.3 Consumption . . . 16
1.4 Portfolio reallocations . . . 17
1.5 Systemic banks . . . 18
1.6 Peer effects . . . 21
1.7 Placebo Test . . . 24
1.8 Observations window robustness check . . . 25

2.1 Descriptive statistics (LLD) . . . 49
2.2 Descriptive statistics (NVM) . . . 50
2.3 The Effect of Leverage Constraints and House Prices on Household Debt . 54
2.4 Bunching at the LTI limit . . . 59
2.5 Bunching at the LTV limit . . . 66

3.1 Descriptive Statistics (CAO Data) . . . 79
3.2 CAO data . . . 81
3.3 Descriptive statistics (MEED) . . . 82
3.4 Descriptive statistics (Consensus data) . . . 85
3.5 Wage Dispersion Test . . . 86
3.6 Extensive margin response to nominal rigidity . . . 89
3.7 Intensive margin response to nominal rigidity . . . 90
3.8 Heterogeneity Analysis . . . 92
3.9 Agreement dates vs Start dates . . . 94

3.10 Robustness: Identification based on Agreement dates 95
3.11 Workers' individual wages 97

4.1 Descriptive statistics 114
4.2 Results: the Netherlands 120
4.3 Results: Italy 121
4.4 Heterogeneity analysis 123

List of Figures

1.1 Empirical distribution of deposits around 38.000 euro 8
1.2 Empirical distribution of deposits around 100.000 euro 8
1.3 Deposited amounts, treatment status 13
1.4 Deposited amounts, treatment status, and the social circle 20
1.5 Graphical test for the parallel trend assumption 23

2.1 Regulatory Loan-to-Income limits (Heatmap). 36
2.2 Changes in maximum allowed indebtedness 37
2.3 (Un)constrained optimal consumption and mortgage distribution. 44
2.4 Exogenous changes in the income classification 51
2.5 Instrumental Variables 53
2.6 Response to a 0.2 LTI limit change 55
2.6 Bunching at the LTV distribution 67
2.7 Pre-treatment LTV distributions 68

3.1 Agreement Delay 80
3.2 Identification Strategy 87
3.3 Extensive and Intensive Margins Response 90
3.4 Heterogeneity in Extensive and Intensive Margins Response 93
3.5 Parallel Trends test 98

4.1 House price trends in Italy and the Netherlands 114
4.2 Expected vs realized house price changes 115
4.3 Distribution of expected house price changes 115
4.4 Maintenance expenses 116

4.5 Correlation between size of home improvements and changes in house values122

Preface

The present thesis contains four essays on *the Micro Effects of Macro Policies*. The aim of the thesis is to understand ultimate consequences for households and firms of policies and institutions that aim at targeting the functioning of entire markets. Most of the focus is on the analysis of macroprudential policies, i.e. policies designed to limit risk intake in the financial system. The prefix *macro* indicates that the policies target the whole or a significant part of the financial system, while *prudential* indicates that the policies aim at ensuring that everyone takes a cautious approach towards potentially systemic risks. Macroprudential policies were introduced within a comprehensive financial assessment framework for banks, pension funds and insurance companies that has deeply impacted the way these institutions operate. The thesis aims to evaluate the effectiveness of these policies and to understand their ultimate consequences on households.

The four chapters have two common and distinctive features. The first distinctive feature is the use of microdata to answer typically macroeconomic questions such as the relation between asset prices and credit, the relation between nominal wage rigidities and labour market fluctuations, the role of public guarantees in enhancing trust, the consumption response to changes in home equity. Thanks to the available data, and thanks to the use of appropriate micro-econometric techniques, the thesis aims at drawing causal conclusions about the micro-effects of the policies and events studied, i.e. the effects on households and firms. The second common feature is that all the chapters focus on the Netherlands, a country which, because of its institutional and socio-economic background, represents an ideal laboratory to address the previously mentioned questions. The thesis in fact pays particular attention to the institutional details and to the policy implications that can be drawn.

In the first chapter, together with Mauro Mastrogiacomo and Giacomo Pasini, I evaluate the effects on household savings of the introduction of the Deposit Guarantee Scheme (DGS). The DGS is a deposit insurance, i.e. a mandatory bank-financed insurance scheme that protects depositors in case of bank default. Deposit insurance schemes prevent one of the major sources of extreme crises in monetary history (bank runs) and they do so without reducing the ability of banks to transform illiquid into liquid assets (Diamond and Dybvig, 1983). For this reason, the DGS is a regulation with important macroprudential and financial stability implications and, along with a common banking supervision and banks resolution scheme, it is one of the three pillars on which the European banking union is build upon. However, during the Global Financial Crisis (GFC) bank runs returned to be a source of public concern, both in the U.S. and Europe. The DGS was introduced in 2008 with the goal of sustaining depositors' confidence towards the banking sector. To achieve this objective, the reform extended the (public) insurance coverage to all deposit accounts up to 100.000 euro. The analysis presented in the first chapter exploits a quasi-natural experiment based on the insurance limit established by the DGS to isolate the confidence effect induced by the reform. Our findings support the use of insurance limit increases as a policy tool to increase confidence in the banking sectors in times of risk and uncertainty.

In the second chapter, I evaluate the effect of macroprudential policies aimed at limiting household indebtedness. A striking empirical regularity of the Great Recession is that the larger the increase in household leverage prior to the recession, the more severe the subsequent recession (Mian and Sufi, 2018). In the U.S., mortgage lending expansion led to sustained increase in house prices (Mian and Sufi, 2009) and household balance sheet effects have accelerated the subsequent drop in spending and consumption (Mian and Sufi, 2010). For this reason, macroprudential tools in the form of Loan to Income (LTI) and Loan to Value (LTV) limits have been introduced to limit risk taking among households, especially during the expansionary phase of the business cycle. The findings shows that, contrarily to what stated by Cerruti and Claessens (2017), there can be important complementarities between LTI and LTV rules. LTV limits alone, despite are found to be the most binding constraint for many households, cannot limit credit growth to households in times of sustained house price increases. Despite the current practice in advanced economies, where either LTI or LTV caps are typically used (Cerruti and Claessens, 2017), our policy implication supports the use of both macroprudential tools, especially in economies with overheating housing markets.

In the third chapter, together with Jante Parlevliet and Mauro Mastrogiacomo, I investigate the consequences of contract staggering in collective agreements. Contract staggering is the result of the institutional environment in which employers and workers bargain wages. In particular, it is a nominal rigidity in wage setting that is typical of centralised barganing systems, and refers to the frequency, the timing and the synchronisation of wage changes. Wages are staggered when contract decisions in the economy are not made at the same time (Taylor, 1979), and are also valid for a number of periods. As a result, only part of the economy can adjust to changes in economic conditions, while the remaining part will be able to adjust once the validity period expires. This property of wage contracts is very well-known in macroeconomics. Models that include such nominal rigidities typically outperform the corresponding models that do not include them (Christiano, Eichenbaum and Evans (2005), Gertler and Trigari (2009)). In these models, nominal rigidities amplify labour market fluctuations, create wage dispersion and aggravate the response of unemployment to a shock. In this chapter we use a quasi-natural experiment based on the staggered start dates of collective agreements to investigate whether nominal wage rigidities, by preventing some firms to adjust contracts to the crisis, led to more layoffs during the 2009 recession. Our findings underscores the role of nominal wage rigidities in the Dutch labour market, as we show that decisions agreed in collective agreements are flexible enough to leave some room for adjustments to the individual employers, not to cause employment reductions.

In the fourth chapter, together with Mauro Mastrogiacomo, I estimate the marginal propensity to consume out of housing wealth. The topic is particularly relevant to understand the transmission of monetary policy on households. In recent years, property valuations have reached new peaks in many countries and cities across the world, and persistent low interest rate environments are contributing to fueling the indebted demand (Mian, Straub, Sufi, 2020) and boosting house prices. The extent to which monetary policies affect the real economy thus depend on the consumption responses of homeowners, who often have a dominant share of their portfolio invested in housing. Also, the topic has important implications also for financial stability: if consumption responses to housing wealth are large, then policies aimed at stabilizing the real estate and the real estate financing market can help achieving levels of consumption smoothing, otherwise impossible to achieve. The paper shows, theoretically and empirically, the role of maintenance and home improvements as an endogenous channel in the estimation of the marginal propensity to consume (MPC) and save (MPS) out of housing wealth. The model shows that, if house qualities are continuous, home improvements are endogenous to the home equity dynamics and, together with consumption and saving, they respond to wealth effects. In the empirical analysis, we propose a comparative empirical approach for the estimation of the MPS, that uses comparable survey data for Italy and the Netherlands. The research question asks whether, consistently with the model, home improvements can bias the estimation of the MPC/MPS. Our results reject this hypothesis and show that, despite a theoretical relevance in considering home improvements, its practical relevance for the estimation of the MPC/MPS is modest. The result suggests that home improvements did not significantly improve the overall quality of properties, and point towards the exogeneity.

Chapter 1

Being in Good Hands: Deposit Insurance and Peers' Financial Sophistication

1.1 Introduction

The ultimate goal of any Deposit Guarantee Scheme (DGS hereafter) is to safeguard the confidence of small savers. Thanks to a deposit insurance that reimburses a limited amount of depositors' savings in case of default, bank runs are prevented, making bank defaults less likely. Bank runs, i.e. mass withdrawals caused by solvency concerns, played a prominent role in monetary history, with recurrent episodes since the early 1900 (Diamond, Dybvig (1983)). In fact, the prevention of bank runs is at the root of deposit insurance.

There is a long-lasting debate on policy interventions aimed at preventing bank runs: Friedman and Schwartz (1963) suggest that payment restrictions such as convertibility

This chapter is a joint work with Giacomo Pasini and Mauro Mastrogiacomo, and has been previously published as a DNB Working Paper. We thank all participants and discussants at seminars at Vrije Universiteit Amsterdam, De Nederlandsche Bank, the 15th Workshop on Social Economy, the International Meeting on Applied Economics and Finance (IMAEF), Ca' Foscari University of Venice, the International Congress in Econometrics and Empirical Economics (ICEEE), St. Gallen University, Univerity of Groningen and the Annual Meeting of the Society of Household Economics for the very useful comments.

suspension ensure reasonably small effects of bank panics on banks balance sheet. Chari and Jagannathan (1988) confirm the result by directly modeling bank runs as situations where agents, observing large withdrawals, correctly infer that the bank is likely to fail and they precipitate to make withdrawals. Diamond and Dybvig (1983) show that a tax-financed deposit insurance dominates convertibility suspension and lending-of-last-resort as alternative solutions.

Despite the long-lasting debate, there is no actual empirical evidence (Iyer, Puri (2012)) and little is known on the effects of deposit insurance. The academic debate concerning deposit insurance has not developed substantially after the key contributions of the eighties, due to a substantial fall in the occurrence of bank runs. However, bank-panics suddenly re-emerged as a source of public concern at the onset of the recent financial crisis (Gertler, Kiyotaki (2015)), and the DGS was introduced with the aim of restoring confidence in the banking sector.

In this paper, we empirically investigate the effect of the Deposit Guarantee Scheme introduction. A unique attempt has been proposed by Iyer et al. (2017) that use data of the population of deposit accounts in Denmark to study the effect of the DGS introduction. According to the authors, the DGS reform in Denmark resulted in an insurance limit decrease, as depositors have always been given implicit government guarantees, and they find increasing withdrawals after the reform. In this paper, we focus on the opposite case of the Netherlands where, as in the majority of EU countries, the reform resulted in a substantial extension of the guarantees provided by the former deposit insurance. Our research question asks whether enlarging public guarantees during crises sustain depositors confidence and, thus, households savings. In other words, we ask whether insurance limit increases can be used as trust-enforcing devices. To address this question, our identification strategy exploits the change in the maximum guaranteed amount to identify the savings response of partially uninsured depositors. The intuition is the following: after the reform all deposits below the 100.000 euro threshold are fully insured, so no reallocation is needed. On the contrary, deposits above the insurance limit are still partially at risk and can be reallocated. If perceived riskiness is high, depositors can reallocate savings to reduce their uninsured balance. Instead, if the insurance limit increase is an effective policy device, it directly affects depositors' perceived riskiness, and reduces their propensity to reallocate deposits as they know that a credible institution is providing direct guarantees. Our findings show increasing savings above the threshold, after the reform, and suggest that an increase in the regulatory limit, by providing additional guarantees, can be an effective trust-enforcing device.

The contribution of this paper is that, unlike previous studies such as Iyer et al. (2017), we take advantage of survey data to test for possible channels explaining our result. We investigate four main channels: consumption responses, flights-to-liquidity, implicit public guarantees, and peer effects. First, the increase in savings may have come at the cost of consumption, especially of durable goods, which is often postponed during crises. Second, while risks in the banking sector were high during the Financial Crisis, risks in financial markets were probably higher, and the increase in savings may be explained by reallocation from risky to (nearly) risk-less assets such as deposits, which are covered by additional guarantees. Third, the increase in savings may reflect implicit other than explicit guarantees, as governments would be called to decide whether to bail out a bank or not, in case of serious financial distress. Fourth, depositors' response can be driven by peer effects: the idea is that when people observe large withdrawals, fears of solvency grow resulting in even larger withdrawals (Chari, Jagannathan (1988)), and models describe withdrawals as strategic decisions, that agents take after observing what the others do. In fact, while bond holders are typically reimbursed on the basis of *seniority* in case of bank default, depositors are reimbursed depending on their "place in line", given the sequential nature of bank refunds (first come, first served).

Our results show that, while the result cannot be explained by consumption cuts or drastic asset re-allocations, savings increased mostly in systemic banks, which have higher public implicit guarantees than small banks, as they are *too big to fail*. Moreover, our results highlight substantial heterogeneity in the response to the reform as a function of the characteristics of depositors' social circles (exogenous peer effects): while people with unsophisticated peers save more as a response to the reform, people with financially sophisticated peers tend to be more cautious and keep their savings close to the insurance limit established by the DGS.

A unique attempt to empirically study peer effects in deposit decisions is due to Iyer et al (2012). Their case study is a run faced by an Indian bank in which, to open a deposit account, one needs an introduction from someone who already has an account opened at the same bank. They show that the probability of running at the bank is increasing in the fraction of people running in the introducer network, and show that a deposit insurance partially helps in mitigating runs. To investigate the effect of peers in deposit and withdrawal decisions, a major challenge is the unavailability of information on network links, such as friendships. This information is available in very few surveys since it is prohibitively expensive to collect. Therefore, peer effects need to be proxied and researchers usually face two possibilities. One is to proxy peer effects with neighborhood effects using location information, under the assumption that social interactions are local. However, also location information is rarely available because of privacy restrictions. Another possibility is to define a social circle on the basis of a common characteristics (such as age and education). In both cases, the econometrician never observes the true links, and the resulting proxy can be very poor. To overcome this issue, we take advantage of the set of Aggregated Relational Data (ARD hereafter, see Breza et al, 2020) available in our survey, i.e. questions of the form: "How many of your acquaintances have *trait k*?" or "Which level of *trait k* do most of your acquaintances have?". The advantage of ARD is that they provide a correct summary statistic of a given social circle's characteristic. By letting the respondent self-define his own social circle, we do not need to proxy his network ties and, thus, its characteristics. In such a way, we solve the issue of measurement error due to the misspecification of the proxy variable. We use ARD to proxy for peers' financial sophistication, using the questions on peers' education and income.

In the literature, there is large empirical evidence that peers influence people's financial decisions such as saving decisions (Duflo, Saez (2012)), stock market participation decisions (Hong et al. (2004), Brown et al. (2008)), borrowing decisions (Haliassos et al. (2014)) and the decision to insure (Cai et al. (2015)). A common explanation is that financial decisions involve complexities that individuals have difficulties in understanding, based on their own education, information and experience (Cai et al. (2015)). Therefore, gathering information is an expensive activity, and peers carry the most informational content. We use ARD on education and income exactly to link between depositors' decisions to the presence of sophisticated peers.

Our main contribution is to show that peer effects and financial sophistication matter even for decisions involving the most basic financial instruments, deposits. It has already been shown that non financial sophisticated households make severe investment mistakes in the stock markets (Calvet et al (2009)). However, while stocks are mainly held by richer and highly educated people (Haliassos, Bertaut (1995), Guiso et al. (2003)), almost everybody hold deposits. Thus, sophistication levels can differ substantially.
The remainder of the paper is organized as follows: in Section II we provide institutional details about the DGS reform and we discuss the framework that motivates our empirical study. Section III discusses data and descriptive evidence. In Section IV we present the empirical analysis on the effect of the reform and on the role of peers. Section V concludes.

1.2 Institutional framework

According to the Financial Stability Forum (2001) the deposit insurance is one of the key elements of the financial safety-net, along with prudential regulation and supervision, and a lender of last resort. By the end of 2008 the EU established the introduction of a new harmonized deposit guarantee scheme able to unify and extend the existing national guarantee schemes. In order not to weigh on taxpayers, the new Deposit Guarantee Scheme is entirely funded by the financial institutions, and reimburses depositors' savings up to 100.000 euro across all EU countries.
Before this reform, deposits in the Netherlands were covered by a national deposit insurance (*Depositogarantiestelsel*) that reimbursed deposits up to 38.000 euro. Therefore, depositors gained additional protection on their savings. More precisely, up to October 2008, deposit insurance in the Netherlands guaranteed a 100% coverage of the first 20.000 euro and an additional 90% coverage of the next 20.000 euro, with a total maximum insurance coverage up to 38.000 euro. The DGS covers private individuals and small businesses and it is limited to bank saving products only. All insurance products such as life insurances are excluded. In particular, the assets covered by the DGS include all payment and saving bank products like payment accounts, demand and fixed-term deposits as well as all credit balances of credit cards.

Another important detail is that the DGS is a per person-per bank insurance scheme, meaning that in case of bank failure the 100.000 euro insurance limit has to be applied to the sum of all amounts of all relevant assets held in the same bank by the same account holder. In other words, while two deposit accounts held in the same bank are covered up to a total guarantee of 100.000 euro, two deposit accounts held in different banks are covered up to a total of 200.000 euro, as the guarantee applies to each bank separately.

1.3 Data and Descriptive Analysis

For the empirical investigation that follows, we use data from the DNB Household Survey (DHS). The DHS is a panel data survey representative of the Dutch speaking population and consists of about 2000 households interviews in every year since 1993.
The DHS contains very detailed information about Dutch households' assets and liabilities, as well as psychological and economic aspects of financial behavior and information about personal characteristics and living conditions. Eventually, DHS contains information both at the household and the household member level.
A nice feature of the DHS panel survey is that respondents are asked to list every single deposit account, as well as any other saving and investment product, and to indicate the corresponding account balance and the financial institution in which the account is registered. For our empirical application we first obtain information on household financial assets by aggregating all assets held by all household members. Then, since the DGS is a per person - per bank deposit insurance that applies to saving products only, we focus on saving deposits and we sum up all deposit amounts held by the same account holder in the same bank, and we reshape the data at the deposit level. After this procedure, we obtain an unbalanced panel where our statistical unit consists of the deposit amount held by each household member in each bank, with one observation for every year the household participates to the survey.
Table 1.1 reports descriptive statistics on saving accounts held by Dutch households. Saving accounts are bank accounts used mainly for saving purposes, in which people usually deposit considerable amounts. These accounts are distinct from checking accounts that

are used mainly for paying or receiving the salary, and that usually have very low deposited amounts. For this reason we exclude them from our analysis. The number of saving deposits that Dutch households hold ranges from one to twenty, even though the majority of them hold one or two saving accounts, which correspond to the median values. The table also shows that most deposits are concentred in the same banks, as the average number of banks is lower than the average number of accounts. The average deposit amount ranges from about EUR 16.000 in 2007 to about EUR 24.000 in 2010. The average deposit amount, as well as the maximum, is quite high and is far above the median value. This statistics excludes saving mortgages, i.e. saving accounts pledge to residential mortgages, in which instead of paying back the principal, the borrower contribute to a saving deposits until maturity, when the amount saved is used to settle the outstanding debt.

Table 1.1: Descriptive statistics

	2007	2008	2009	2010
Deposit amount				
mean	16.717,88	16.946,81	18.865,74	24.338,34
median	5.597,00	6.674,00	7.966,50	9.000,00
max	556.586,00	650.759,00	850.000,00	500.000,00
N. of accounts				
mean	1.16	1.24	1.32	1.34
median	1	1	2	2
max	15	20	20	17
N. of banks				
mean	1.15	1.17	1.23	1.25
median	1	1	1	1
max	4	4	4	5

Note: Descriptive statistics at household level. The top Panel the mean, the median and the maximum amount deposited in a saving account by the Dutch households in each year. The middle panel reports the mean, the median and the maximum number of saving accounts that Dutch households hold in each year. The bottom panel reports the mean, the median and the maximum number of different banks in which Dutch households hold saving accounts, for each year.

Figure 1.1 and 1.2 show the distribution of deposit amounts before and after the policy change, in a narrow window around the old and the new insurance limit, respectively. The pre reform period is 2007-2008, the post reform period 2008-2009. The old insurance limit is 38.000 euro, the new limit established by the DGS is 100.000 euro. In both figures, the vertical line denotes the insurance limit.

Figure 1.1: Empirical distribution of deposits around 38.000 euro

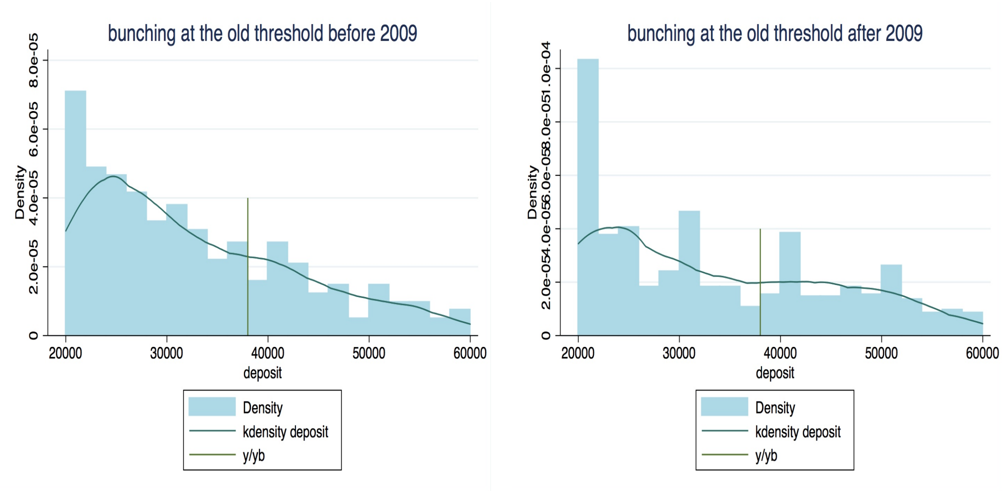

(a) (b)

Note: The figure shows the distribution of deposit amounts around the old EUR 38.000 before (a) and after (b) the DGS reform. The vertical line denotes the EUR 38.000 insurance limit.

Figure 1.2: Empirical distribution of deposits around 100.000 euro

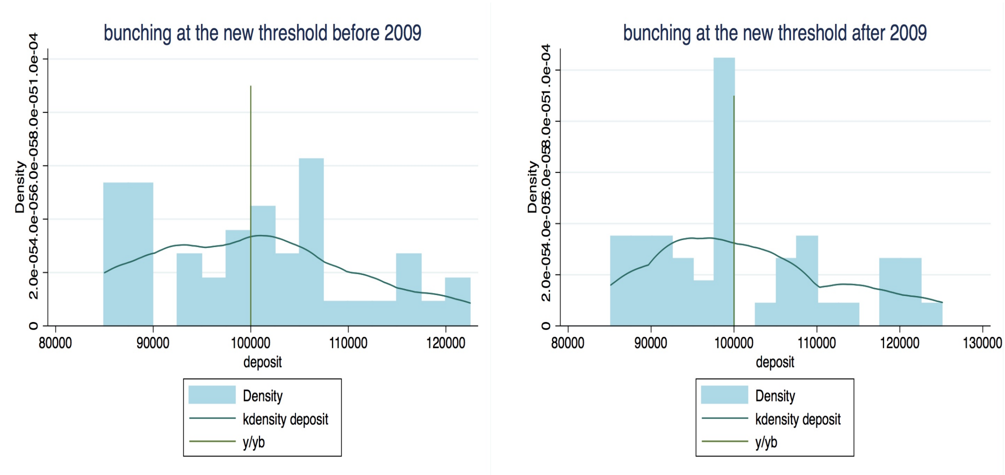

(a) (b)

Note: The figure shows the distribution of deposit amounts around the new EUR 100.000 before (a) and after (b) the DGS reform. The vertical line denotes the EUR 38.000 insurance limit.

In Figure 1.1, the deposit mass just below EUR 38.000 in panel (a) disappears after the policy change in panel (b). Possibly, part of the deposits mass below 38.000 euro reflect strategic allocations (bunching at the old insurance limit). This is confirmed in Figure 1.2, showing the same comparison (the amount distribution before and after the policy change) around the new insurance limit. A higher mass of deposits around EUR 100.000 emerges, with a density that is two or three times higher than the density in nearby bins. Moreover, there is no mass in correspondence of deposit amounts just greater than 100.000 euro. Figure 1.2 provides descriptive evidence that is suggesting that some depositors (*i*) strategically bunch at the threshold and (*ii*) adjust their allocations, either because of changes in perceived risk or changes in the amount of insurance. This evidence, being descriptive, cannot lead us to conclude that there have been *silent bank runs* after the reform, i.e. withdrawals up to the insurance limit. However, it suggests that that the insurance limit is a reference point for some depositors, and motivates the empirical analysis presented in the next section.

1.4 Empirical Analysis

1.4.1 Wealth allocation under the DGS rules

Suppose an agent is endowed with a liquid amount of wealth a_t and uses bank deposits as saving device. Banks offer homogeneous saving products, that is they all offer a deposit contract that gives at time $t+1$ a fixed rate of return $R = 1 + r$ for each unit of saving deposited at time t. Next, suppose that a DGS is in place, so that deposits are covered by a deposit insurance up to a threshold τ. As a consequence, for a depositors endowed with $a_t \leq \tau$, the evolution of wealth is deterministic and equal to $a_{t+1} = Ra_t$. Next, assume that a positive bank default probability π exists, so that depositors endowed with a wealth amount higher than the insurance threshold ($a_t > \tau$) risk losses on their deposits.
In this case the the expected saving amount in the next period is equal to:

$$E(a_{t+1}) = R\tau + [\pi\lambda + (1-\pi)R](a_t - \tau) \tag{1.1}$$

Where λ is a recovery rate. With probability π the bank fails and returns to the depositors the deposit insurance amount τ, plus the recovery rate $\lambda \in (0,1)$ on the uninsured balance. With probability $(1-\pi)$ the bank doesn't fail and the agent receives the agreed rate of return $R = 1 + r$ on the whole wealth. Therefore, the expected future level of wealth is made up of a $safe$ part equal to $R\tau$ and a *risky* part that is equal to $[\pi\lambda+(1-\pi)R](a_t-\tau)$.

Next, a particular feature of the DGS is that it is a *per person*, *per bank* deposit insurance, meaning that depositors are insured up to τ in each different bank they are saving in. This means that for agents endowed with $a_t > \tau$ nothing is lost: they can open deposit accounts in different banks and keep their wealth fully secured by saving an amount at most equal to τ in each account opened. However, opening several deposit accounts is not for free: it typically requires a fixed transaction cost (service fees, taxes) to be paid, other than a general implicit cost (inconveniences, such as many passwords, notifications, bank cards renewals etc.) involved by holding multiple deposits. If the former cost increases linearly with the number of accounts being held, the latter can possibly increase more than proportionally relative to the number of accounts. As a result, a depositor endowed with $a_t = \tau$ in $n-1$ saving accounts, but leaving part of his wealth unsecured on his n-th saving account, will have an expected wealth dynamics equal to:

$$E(a_{t+1}) = Rn\tau + [\pi_i\lambda + (1-\pi_i)R](a_t - n\tau) - c(n) \tag{1.2}$$

Where $c(n)$ is a cost function, with $c'(n) > 0$, and π_i is now the probability of default associated to bank i where the depositor holds a partially uninsured deposit account, i.e. an amount greater than τ[1]. As it can be easily seen, the wealth dynamic has a cost-benefit structure in which the benefit is represented by the insurance coverage that is increasing and linear in n. The cost of holding is also increasing in n, possibly more than proportionally. As a result we can write the expected wealth dynamics for an agent having $i = 1, .., n$ deposit accounts as follows:

$$E(a_{t+1}) = \sum_i \underbrace{min\{Ra_{i,t}\ ;\ R\tau\}}_{\text{insured amount}} + \underbrace{max\{0\ ;\ [\pi_i\lambda + (1-\pi_i)R](a_{i,t} - \tau)\}}_{\text{uninsured amount}} - c(n) \tag{1.3}$$

How do perceived riskiness and deposit insurance affect households savings decisions? Specifically, what is the response to changes in the insurance limit or to changes in perceived riskiness?

[1] In fact, in case of multiple accounts holding, each account is fully insured if and only if $\omega_j a_t \leq \tau\ \forall j$, where ω_j is the fraction of wealth deposited in the $j-$th account. If all deposit accounts are insured, the overall wealth is fully insured too.

Eq. 1.3 show that while an increase in the deposit insurance limit τ affect everybody, as one either gains full insurance coverage or a reduction of the uninsured balance, increases in banks' default probability π only affect those having partially uninsured accounts in those banks, i.e. deposits above the limit. The next section describes the set up of a quasi-natural experiment based on this intuition, that we use to study how changes in deposit insurance affect household savings when they are exposed a bank default risk.

1.4.2 Identification strategy

In this section we analyze the households' response to the 2008 DGS reform to see whether insurance limit increases sustain depositors' confidence. To target depositors having incentives to reallocate their deposits, we take the new DGS insurance limit as a sharp rule to define the treatment (T) and the control (C) of a Difference-in-Differences (DiD) experiment. We set as T units all deposit accounts with an amount greater than the insurance limit threshold: $T = 1[a_t > \tau]$. Conversely, the C group consists of deposits with amount below the limit. Again, after the reform all deposits below the 100.000 euro threshold are fully insured, so no reallocation is needed. On the contrary, deposits above the insurance limit are still partially at risk and can be reallocated. Therefore, risk exposure represents the treatment of our quasi-natural experiment: if riskiness increases depositors should withdraw their uninsured balance; if instead perceived riskiness decreases (thanks to the additional guarantees provided), depositors will be more willing to let uninsured amounts on their saving accounts.

Our DiD compares T and C outcomes in a sufficiently narrow window of observations around the insurance limit, in the years before and after reform. As a result, the resulting estimate will reflect a local effect. The choice of the observation interval poses a tradeoff: on one hand, by excluding T and C units too far from the threshold we can identify the local effect properly, but we lose estimation precision as the number of observations falls. On the other hand, taking wider intervals increases the sample size and the estimation accuracy, but we would include depositors that, being too far away from the insurance limit, may not respond to its change. According to this tradeoff we select the $(20.000 - 300.000)$ interval. Also, we select a four year window around the reform, from 2007 to 2010. Given these premises, we consider the following linear model for DiD:

$$y_{i,j,t} = c_i + \lambda_t + \beta\, T_{i,j,t} + \sum_{t=1}^{T} \gamma_t^{ATT} \big(T_{i,j,t} \times \mathbf{1}[t \geq \tau]\big) + \mathbf{X}_{j,t}\delta + \epsilon_{i,j,t} \tag{1.4}$$

The dependent variable is the amount deposited in account i by household j at time t and $T_{i,j,t}$ is the treatment group identifier. The average treatment effect on the treated (ATT) is captured by the coefficients $\gamma_t^{ATT} = E(y_{\tau+t}(1) - y_{\tau+t}(0)|T = 1)\ \forall t$.

Eventually, $\mathbf{X}_{j,t}$ are household level covariates and c_i are unobserved common factors at the deposit level[2]. The selected covariates are: household income, the outstanding total debt towards the bank where the deposit is registered, a dummy variable indicating whether this year expenses are going to be unusually high and a set of demographics, a quadratic in age and employment status dummies.

Moreover, following Iyer and Puri (2012) we account for bank's cross selling with a variable ranging from 1 to 6 indicating the number of different contract types held with the same bank, including deposit and checking accounts, mortgages and personal loans, deposit books and saving certificates[3]. In fact, the higher is the number of contracts with a customer, the higher is the amount of soft information that the bank obtains, and the higher will be the opportunity cost of changing the bank, since it is unlikely for the customer to find the same economic conditions in other banks. For this reason, they claim that cross selling represents for banks a complementary insurance against the risk of having a run. The next section discusses the estimation results.

[2]Note that due to our selection, the great majority of households has either zero or one account included in our estimating sample. Therefore, c_i corresponds to a common factor at household level. For the few households with more than one account in the interval EUR 20.000-300.000 instead, c_i allows for a deposit-specific intercept, but we account for the correlation between deposits held by the same household by clustering the standard error at the household level. Even though we think that the relevant heterogeneity to be accounted for is at the household level (risk-aversion), we think that the cost in terms of assumption is lower than the gain in flexibility: assuming a deposit-specific intercept allows us to do within estimation, instead of estimating a version of eq. (4) augmented with all time-invariant controls and all household fixed-effects, which suffers from curse of dimensionality.

[3]Iyer and Puri (2012) can only use the loan link indicator as a proxy for the depth of bank-client relationship (cross selling), i.e. whether the depositor has a loan with the same bank.

1.4.3 The effect of the DGS reform

This section presents the result of the effect on savings of the introduction of the DGS reform. Table 1.2 reports the results obtained from the estimation of eq. 1.4. Standard errors are clustered at household level. Specification (a) is without controls, in (b) we only control for basic demographics and in (c) we use the full set of covariates listed in the previous paragraph. Table 1.2 shows a positive effect on saving amounts of the DGS reform: deposits above the threshold significantly increase in 2009, relatively to the corresponding control-group trend, just a few months after the reform took place (October 2008). The average treatment effect on the treated (ATT), as captured by the coefficient of $above \times 2009$, is positive and statistically significant and its magnitude indicates that in 2009 the average deposit amount above the insurance limit increased on average by EUR 11.720,00 relative to the pre-reform trend[4]. Moreover, the magnitude and the significance of the ATT is robust across all three specifications.

Figure 1.3: Deposited amounts, treatment status

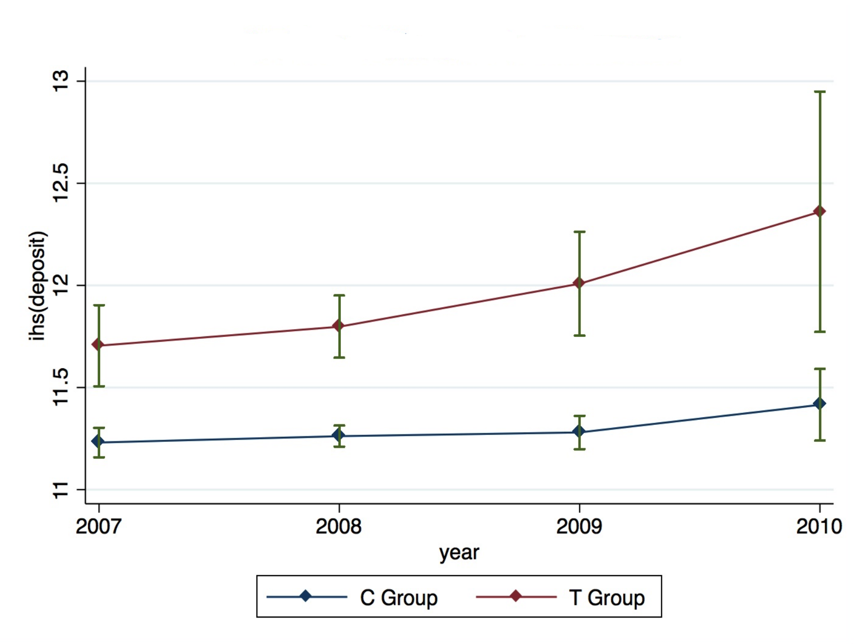

Note: The figure the trends of the amounts deposited in saving accounts above (red) and below (blue) the DGS threshold. The amounts are expressed in logarithms.

[4]See Appendix 2 for the estimation of the ATT magnitude.

Table 1.2: The general effect of the DGS reform

	Dep. variable: Saving amount		
	(a)	(b)	(c)
above	0.489***	0.486***	0.444***
	(0.0936)	(0.0950)	(0.1201)
2009	0.031	-0.032	-0.023
	(0.0354)	(0.0461)	(0.0483)
2010	0.167**	0.085	0.128
	(0.0844)	(0.0979)	(0.0984)
above×2009	0.237*	0.248*	0.311**
	(0.1233)	(0.1264)	(0.1420)
above×2010	0.451	0.455*	0.454
	(0.2991)	(0.2739)	(0.2888)
household income ('000s)			0.002**
			(0.0012)
outstanding bank debt			0.001
			(0.0008)
high future expenses			0.041
			(0.0569)
cross selling			0.152
			(0.0992)
controls	NO	YES	YES
R-squared	0.221	0.235	0.261
N obs. (Nt)	1134	1122	1023
N deposits (N)	517	509	452
N above	77	77	70

Note: The dependent variable is the amount deposited in the saving account i, by household j at time t,expressed in logs. All estimates are Fixed-Effects Difference-in-differences. Standard errors are clustered at household level. The set of covariates include the net annual household income, the outstanding debt towards the same bank where the account is registered, a dummy equal to one when the household head expects high household expenses over the year, and a variable equal to the number of different contract types the household has in the bank. The set of controls include age, its square, and a set of employment status dummies. The symbols *, **, and *** denote ten, five and one percent statistical significance levels, respectively.

Figure 1.3 provides a graphical representation of the result in Table 1.2, showing the average saving amounts above and below the threshold, in each of the sample years. It shows that the trends in the two groups are parallel in the pre-reform period, while between 2008 and 2009 (as the reform is introduced) the average amount in treated deposits increases sharply relative to the average amount in the control group, which remains stable. Then, the two trends seem to return parallel in 2010.

While deposits below the EUR 100.000 threshold are fully secured after the DGS reform, deposits above the threshold are still subject to bank default risk. On one side the introduction of the reform, together with the conditions of the banking sector, reveal the presence of increasing risks on the banking sector. On the other side, the introduction of the DGS reform was exactly aimed at protect depositors and increasing their confidence towards the banking, by signaling that the EU is willing to insure depositors' losses in case of bank defaults. As a result, the perception of increased riskiness can induce depositors to withdraw uninsured balance amounts, while increasing public guarantees, if effective, can lower depositors' propensity to withdraw their savings. Based on this evidence reported in Table 1.2 we conclude that the DGS reform was effective in increasing depositors' confidence, as depositors (still) exposed to banks' default risk have increased, on average, the amount deposited in their savings. The next four sections investigate possible additional explanations for this result.

1.4.4 Consumption cuts

Saving decisions are thought as simultaneous to consumption decisions. Therefore, the increase in savings after the reform can possibly be explained by consumption cuts. During the crisis, households may face increasing future income risks, and this can induce to precautionary save (Skinner, 1988). Of course, durable consumption goods are easier to postpone than necessary goods. Here, we move from a deposit level analysis to a household level analysis and we set as treated units the households having at least one uninsured deposit account, in order to have a one-to-one correspondence with the deposit level analysis. Our dependent variable is a dummy equal to one when the respondent is planning to do big expenses in durable consumption (such as car or furniture purchases) during the year. Again, to be consistent with the previous analysis, we exclude families having all their deposits outside the EUR 20.000-300.000 interval. Results are reported in Table 1.3 and show that consumption cuts were not responsible for the increase in savings: the mild significance of $above \times 2010$ vanishes as we control for our set or covariates.

Table 1.3: Consumption

	Dep.variable: Durable consumption		
	(a)	(b)	(c)
Above	0.168	0.164	0.166
	(0.1330)	(0.1433)	(0.1477)
2009	-0.037	0.017	0.029
	(0.0378)	(0.0651)	(0.0665)
2010	-0.110	-0.025	-0.020
	(0.0871)	(0.1240)	(0.1277)
Above×2009	-0.273	-0.272	-0.277
	(0.1796)	(0.1878)	(0.1752)
Above×2010	-0.017	-0.001	-0.007
	(0.1246)	(0.1335)	(0.1353)
controls	NO	YES	YES
N° obs.	911	911	864

Note: The dependent variable is a dummy variable equal to one when the respondent has plans to do high expenses in durable consumption during the year. All specifications are at the household levels. All estimates are Fixed-Effects Difference-in-differences. The symbols *, **, and *** denote ten, five and one percent statistical significance levels, respectively.

1.4.5 Portfolio reallocations

The increase in savings above the threshold is due portfolio reallocation of investors. If the banking sector was in distress due to the economic crisis, the situation wasn't any better in financial markets. As an example, the Dutch stock market index AEX lost nearly 50% of its market capitalization during 2008. The increase in savings above the threshold after the reform can be potentially explained by portfolio reallocations, i.e. investors converting their risky assets in liquid assets.

Here, we regress the same DiD specification, again at the household level, using as dependent variable the amount of financial wealth, defined as any asset other than those covered by the DGS (saving deposits and saving certificates), business equity and housing. It includes life insurances, bond and stock holdings, and shares of mutual funds. If investors convert stocks into liquidity, which they pledge in a deposit account, we would observe a decrease in the value of financial assets, among those having deposit accounts above the insurance limit. Table 1.4 shows that even though the ATT is negative and significant in the first specifications, once again the significance vanishes as we estimate our most complete specification (c), that includes all control variables. The result lead us to conclude that the increase in savings was not driven by a *flight-to-liquidity* mechanism, i.e. reallocation from risky to risk-less assets among depositors having account balances above the limit.

Table 1.4: Portfolio reallocations

	Dep. variable: Financial Assets		
	(a)	(b)	(c)
Above	-0.095	-0.144	-0.039
	(0.2245)	(0.2226)	(0.2385)
2009	-0.355**	-0.068	-0.091
	(0.1612)	(0.1953)	(0.1928)
2010	-0.616*	-0.073	-0.096
	(0.3511)	(0.4096)	(0.4226)
Above×2009	0.000	0.078	0.046
	(0.4096)	(0.4047)	(0.4141)
Above×2010	-0.902	-0.786	-0.777
	(0.7320)	(0.7687)	(0.7472)
controls	NO	YES	YES
N	965	965	864

Note: The dependent variable is the amount of financial assets (life insurances, bonds, stocks, shares of mutual funds), expressed in logs. All specifications are at the household level. All estimates are Fixed-Effects Difference-in-differences. The symbols *, **, and *** denote ten, five and one percent statistical significance levels, respectively.

1.4.6 Implicit public guarantees

We investigate the role of implicit guarantees by re-estimating eq. 1.4 on the subsample of deposits held in systemic banks. By comparing the estimated ATTs with those obtained from the baseline estimation, we can assess whether the increase in deposits above the threshold after the DGS reform was driven by increasing deposits in systemic banks.
In fact, a possibly unintended consequence of the DGS is to contribute to level the playing field in the funding competition among banks, as deposits below the threshold in systemic or non-systemic banks are equally safe. On the other side, deposits above the threshold held in systemic banks should be considered relatively safer, as these banks have addi-

tional regulatory requirements, as well as implicit public guarantees of bailouts in case of financial distress. In fact, systemic banks are also perceived to be *too big to fail*, due to their size and their interconnections. Table 1.5 shows the result of this subsample analysis. Since the official definition of systemically important financial institutions (SIFI) was provided by the Financial Stability Board only in 2011, we define systemic banks based on market shares. Systemic banks are the three largest banks: Abn Amro, Rabobank and Postbank. The three specifications corresponds to those of Table 1.2.
Results suggests that most of the increase in savings occurred in systemic banks. In fact, the size of the estimated ATTs (in line with those in the baseline estimate) together with the high market shares of these banks (about 80% of respondents' deposits are in these banks), lead us to conclude that most of the increase in savings are in accounts registered in systemic banks. This is consistent with the fact that, due to implicit public guarantees, holding partially uninsured accounts in systemic banks can always be considered safer than those in non-systemic banks.

Table 1.5: Systemic banks

	Dep. variable: Saving amount		
	(a)	(b)	(c)
above	0.552***	0.544***	0.487***
	(0.1108)	(0.1161)	(0.1438)
2009	-0.017	-0.102*	-0.077
	(0.0460)	(0.0557)	(0.0610)
2010	0.118	-0.001	0.059
	(0.0892)	(0.1034)	(0.1072)
above×2009	0.257**	0.275**	0.342**
	(0.1146)	(0.1204)	(0.1465)
above×2010	0.442	0.454*	0.474*
	(0.2870)	(0.2749)	(0.2872)
controls	NO	YES	YES
N obs.	726	721	660

Note: The dependent variable is the amount deposited in the saving account i, by household j at time t, expressed in logs. All estimates are Fixed-Effects Difference-in-differences. All specifications are at the deposit level. Standard errors are clustered at household level. The symbols *, **, and *** denote ten, five and one percent statistical significance levels, respectively.

1.4.7 Peer effects

We eventually investigate whether the results can be explained by peer effects. There are at least three motivations for looking at peer effects in deposits and withdrawals decisions is threefold. The first reason is that there is huge evidence in the Household Finance literature that peers influence any people's financial decisions. The common explanation is that gathering information is a difficult and expensive activity, so people often rely on observational learning and peers' advice. Second, the theoretical literature looks at bank runs (i.e. mass withdrawals) as decisions that agents take after observing what the others do: if agents observe long queues at the bank, they correctly infer that the bank is likely to fail and they also precipitate to make withdrawals, given the sequential nature of bank reimbursement. The third reason is that in the period we consider, a bank run to a Dutch bank was caused by a similar episode: a representative of the unsatisfied bank customers publicly motivated all depositors to withdraw funds from the bank, and his suggestion resulted in a bank run (DNB Annual report, 2009).

Given the absence of information on peers' savings, and thus on their deposits and withdrawals, we follow Calvet et al (2009) who look and the relation between household characteristics and investment mistakes, and show that financial unsophisticated households make severe investment mistakes[5]. Here instead, we look at the relation between peer characteristics and individuals' saving decisions (exogenous peer effects) to see whether depositors with financially sophisticated peers respond differently to the DGS reform.

We employ a Triple Differences (TD) estimator that enables us to investigate heterogeneous effects, i.e. within group differences in the response to the DGS reform, while preserving the identification strategy being used, as the underlying assumptions are almost equivalent[6]. More precisely, within the group of depositors with uninsured balance accounts after the reform, the TD compares the post-reform outcomes of households linked to richer and more educated peers. We proxy peers' financial sophistication with peers' income and education. We first define the degree of peers' financial sophistication as the

[5]In this literature, financial sophistication is proxied by income, wealth and education. Related papers include Calvet et al. (2007), Yannis et al. (2008), VIssing-Jorgensen (2003), Ravi and Zhu (2006).

[6]Lee (2005) shows that once the parallel trend assumption holds, the corresponding assumption for TD. The assumption frequires the difference between the T and C time effects to be the same for both G subgroups.

level of income and education that most people in each respondent's social circle have, by using the corresponding set of Aggregated Relational Data (ARD), and then we build a binary indicator that measures peers' financial sophistication on a relative basis, and we define respondents having sophisticated friends as those having friends that are richer and more educated than they are[7]. Let $G_j = \{0, 1\}$ be the indicator for respondents having sophisticated peers. The TD estimator identifies the following statistic:

$$TD = E\big(y_{\tau+t}(1) - y_{\tau+t}(0)|X, G = 1, T = 1\big) \qquad \forall t \tag{1.5}$$

The TD estimator identifies the average treatment effect on the subgroup G of the treatment group T. In our case, it identifies the effect of the DGS on respondents having uninsured balance accounts ($T = 1$) and financially sophisticated friends ($G = 1$).
Results of a linear model for TD are reported in Table 1.6 and a graphical representation of the result is given in Figure 1.4.

Figure 1.4: Deposited amounts, treatment status, and the social circle

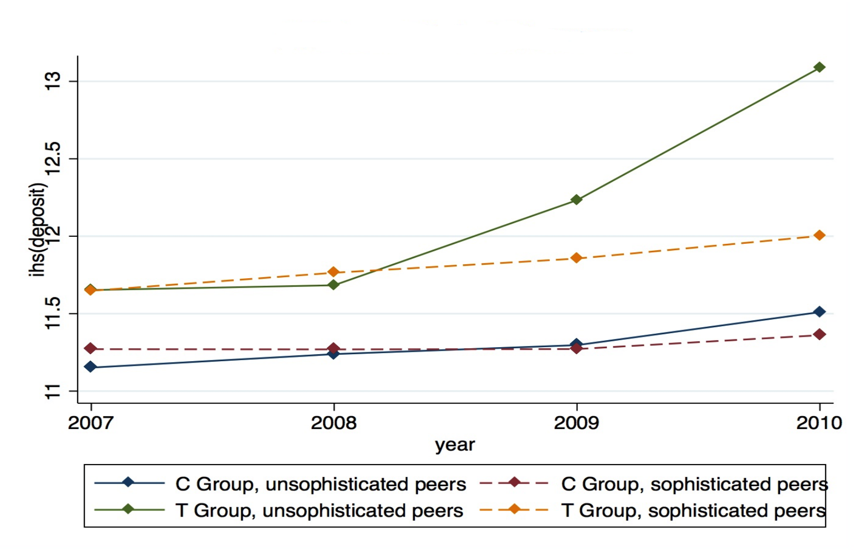

Note: The figure the trends of the amounts deposited in saving accounts above and below the DGS threshold, by depositors with (dashed lines) and without (solid lines) financially sophisticated peers. The amounts are expressed in logarithms.

[7]Of course, we also control for own financial sophistication, namely own income and education.

Table 1.6: Peer effects

	Dep. variable: Saving amount		
	(a)	(b)	(c)
above	0.485***	0.485***	0.410***
	(0.1145)	(0.1024)	(0.1371)
2009	0.102*	0.032	0.040
	(0.0612)	(0.0723)	(0.0715)
2010	0.306**	0.228	0.275*
	(0.1279)	(0.1494)	(0.1511)
soph.peers	0.082*	0.061	0.036
	(0.0442)	(0.0465)	(0.0495)
above×2009	0.442**	0.461**	0.550**
	(0.1973)	(0.1890)	(0.2196)
above×2010	1.080***	1.019***	1.070***
	(0.1846)	(0.1932)	(0.2028)
soph.peers×above	-0.074	-0.061	-0.014
	(0.1256)	(0.1198)	(0.1349)
soph.peers×2009	-0.104	-0.094	-0.098
	(0.0758)	(0.0766)	(0.0736)
soph.peers×2010	-0.217	-0.232	-0.246
	(0.1644)	(0.1701)	(0.1674)
soph.peers×above×2009	-0.263	-0.285	-0.327
	(0.2315)	(0.2235)	(0.2348)
soph.peers×above×2010	-0.856***	-0.778***	-0.854***
	(0.2339)	(0.2559)	(0.2551)
Controls	NO	YES	YES
N obs.	1134	1122	1023
N deposits	517	509	452

Note: The dependent variable is the amount deposited in the saving account i, by household j at time t, expressed in logs. All estimates are Fixed-Effects Triple-differences. Standard errors are clustered at household level. The symbols *, **, and *** denote ten, five and one percent statistical significance levels, respectively.

The within group ATTs in 2009 are identified via the coefficients of *Above×2009* and of *soph.peers×Above×2009*. The first coefficient is significantly positive in all specifications while the second is not. This means that, the increasing savings above the threshold after the DGS reform, is entirely due to depositors with unosphisticated peers, right after the reform. In 2010 the difference in saving behavior widens: the coefficient of *soph.peers×Above×2010* is negative and statistically significant, and fully compensates the size the positive coefficient of *Above×2010*. Relatively to depositors without sophisticated peers, the change in deposit amounts above the threshold of depositors with sophisticated friends is equal to -13.000 euro in 2009 and -36.000,00 euro in 2010[8]. In words, this result shows that among those having partially uninsured accounts (treated depositors), respondents without sophisticated peers save more after the reform, while respondents with sophisticated peers stay more cautious: they do not increase their uninsured balances in a period of unusual uncertainty, and they they keep their deposit amount closer to the insurance limit. Also it shows how peer influence and financial sophistication can have a role also for the decisions regarding the simplest financial instruments. Figure 1.4 shows that before the reform respondents with (dashed lines) and without (solid line) sophisticated peers share the same trends, while after the reform the increase in savings is entirely driven by the latter group. Our interpretation is that, after the bankruptcy of DSB bank in 2008, financially sophisticated depositors may have realized that bank bailouts cannot be given for granted, as stated by the Dutch National Bank (DNB, 2009).

1.5 Internal validity and robustness

Our identification strategy relies on three main assumptions. First, Difference-in-differences and Triple-differences identify the treatment effect on the treated under the assumption that, in absence of the treatment, the T and C groups would follow parallel trends. Second, in our design we assign deposits to treatment and control group on the basis of risk exposure, i.e. whether depositors have uninsured balances (deposits above EUR 100.000) after the reform. Third, we claim that the effect is local and the results in Table 1.2 are estimated using a specific window of observations around the new EUR 100.000 limit.
To test the validity of the common trend assumption, we exploit the panel feature of our data and we use the four years period prior to the reform to repeat specification (c) of Table 1.2. Correspondingly, we pretend the reform year to be the second, 2006. Results are reported in Figure 1.4: over the period 2005-2008 the T group (red solid line) and the C group (blue dashed line) shared the same untreated response, since the two lines are perfectly parallel. This provides further evidence of the plausibility of the common trend assumption and of the validity of our identification strategy.

[8] See again Appendix 2 for the estimation of the ATT magnitude.

Figure 1.5: Graphical test for the parallel trend assumption

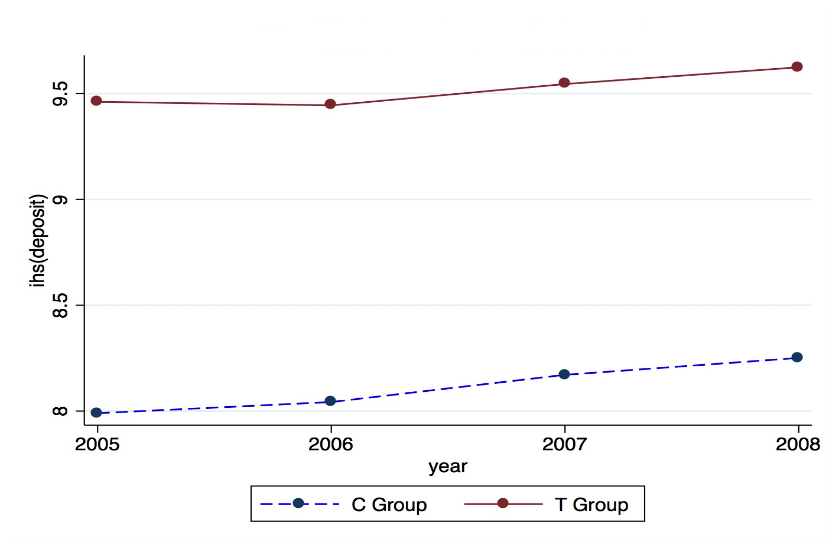

Note: The Figure shows a graphical test for the parallel trend assumption of Difference-in-Differences. It shows the trends of the amounts deposited in saving accounts above (red) and below (blue) the DGS threshold, in the four years period before the DGS reform. The amounts are expressed in logarithms.

Second, in our design we assign deposits below EUR 100.000 to the control group, since they are subject to full insurance coverage. Conversely, deposits above EUR 100.000 are assigned to the treatment group, as they are still at risk and can possibly be reallocated. As a consequence, after the reform the dynamics of the latter should reflect perceived risk, other than household specific needs of consumption and saving. To validate our identification, we perform the following Placebo test: we set as (placebo) treatment group all deposits *below* 38.000 euro. These deposits are subject to full insurance coverage before and after the DGS reform, so the reform didn't affect them as they have been never exposed to default risk. Then, we repeat the same DiD specification of Table 1.2, using now as treatment the group that is known not to be affected by the policy (Placebo treatment). Results are reported in Table 1.5[9]: across all specifications neither the coefficient of $below \times 2009$ nor the coefficient of $below \times 2010$ are statistically significant. A small level of significance emerges in 2010 only, but the effect suddenly vanishes as we control for our set of covariates. Again, this result seems to validate our empirical design.

[9]Since we set all deposits below 38.000 as (placebo) T group, differently from our baseline specifications, we also consider all deposits below 20.000 euro in the estimate, in order to preserve balance across the T and C group. For such a reason, the number of observations is now higher than before.

Table 1.7: Placebo Test

	Dep. variable: Saving amount		
	(1)	(2)	(3)
below	-1.535***	-1.513***	-1.551***
	(0.1532)	(0.1484)	(0.1640)
2009	0.067	0.170	0.262**
	(0.0969)	(0.1081)	(0.1172)
2010	0.194	0.381**	0.495**
	(0.1385)	(0.1891)	(0.1947)
below×2009	0.012	0.007	-0.044
	(0.1166)	(0.1202)	(0.1246)
below×2010	-0.407	-0.459*	-0.241
	(0.2581)	(0.2357)	(0.1999)
Controls	NO	YES	YES
N obs.	4820	4719	4094

Note: The dependent variable is the amount deposited in the saving account i, by household j at time t, expressed in logs. All estimates are Fixed-Effects Difference-in-differences. Standard errors are clustered at household level. The Table reports the coefficients associated with the ATT of interest. *below* denotes the Placebo Treatment. The symbols *, **, and *** denote ten, five and one percent statistical significance levels, respectively.

Eventually, we investigate whether results are driven by a specific selection of the observations window. On the one hand by taking tighter intervals we carefully target the local effect, but on the other hand estimation precision falls because of the loss in the number of observations, especially in the left tail of the deposit distribution. The interval chosen for the baseline regressions was the (20.000-300.000) window. Here we check whether our baseline results of Table 1.2 change as the observations window changes. The results obtained using specification (c) of Table 1.2 are reported in Table 1.6 and show that as we further restrict the observation interval, the ATT remain robust and stable across all different selections.

Taken as a whole, these robustness checks allow us to conclude that the identification strategy seems to be valid since (i) the estimates are able to detect the true null hypothesis of no treatment effect when the given treatment is placebo and (ii) the estimates are not driven by specific window selections.

Table 1.8: Observations window robustness check

	Dep. variable: Saving amount				
amounts window ('000)	18-350	20-280	20-320	22-280	25-250
above	0.502***	0.468***	0.444***	0.458***	0.453***
	(0.1311)	(0.1255)	(0.1201)	(0.1155)	(0.1169)
2009	-0.088*	-0.032	-0.023	0.046	0.064
	(0.0503)	(0.0481)	(0.0483)	(0.0512)	(0.0476)
2010	0.086	0.110	0.128	0.141	0.147
	(0.0923)	(0.0988)	(0.0984)	(0.1114)	(0.1032)
above×2009	0.467**	0.303**	0.311**	0.268*	0.266*
	(0.1830)	(0.1439)	(0.1420)	(0.1362)	(0.1373)
above×2010	0.592*	0.623*	0.454	0.481*	0.492*
	(0.3252)	(0.3185)	(0.2888)	(0.2891)	(0.2777)
Controls	YES	YES	YES	YES	YES
N obs.	1022	1028	931	817	1141

Note: The dependent variable is the amount deposited in the saving account i, by household j at time t, expressed in logs. All estimates are Fixed-Effects Difference-in-differences. Standard errors are clustered at household level. The Table reports the coefficients associated with the ATT of interest. Each column reports a different observation window selection, from the most wide to the most narrow. The symbols *, **, and *** denote ten, five and one percent statistical significance levels, respectively.

1.6 Concluding Remarks

This paper evaluates the effect of the introduction of the Deposit Guarantee Scheme (DGS). The DGS was introduced by the EU to extend and harmonize all national deposit insurance schemes in Europe, and its goal was of sustain depositors' confidence in a period of unusual uncertainty for the European banking sector. In this paper we look at the particular case of the Netherlands, where the reform resulted in an insurance limit increase during a period characterized by banks' defaults (DSB, Icesave) and bailouts (Fortis).

To isolate the confidence effect, we look at the post-reform changes in savings of deposit accounts above the insurance limit established by the DGS. This insurance limit distinguishes insured deposits from (partially) uninsured ones in the distribution of savings. In particular, using Difference-in-Differences we compare insured and partially uninsured deposit accounts to disentangle the effect that larger public guarantees have on perceived riskiness. We find significant effects of the DGS reform on saving amounts: results show an average increase of uninsured saving amounts by EUR 11.700 relative to the pre-reform trend. We conclude that insurance limit increases are effective policy tools in boosting depositors' confidence, as they signal an institutional interest in avoiding depositors' losses. We document that the increase in savings above the insurance limit occurred mostly in systemic banks, and wasn't explained neither by consumption cuts nor by flight-to-liquidity behaviors, i.e. conversions of risky assets into risk-less and liquid assets. Moreover, results show households' response heterogeneity to the DGS reform as a function of the characteristics of people's social environment (exogenous peer effects): among depositors with partially uninsured accounts, we find that people with unsophisticated peers save more as a response to the reform, while people with sophisticated peers save more cautiously, keeping their savings closer to the insurance limit established by the DGS. The Triple Difference estimate indicates a difference up to EUR -36.500 among the two groups of treated households. We argue that the possible channel from peers' sophistication to individual savings decisions is concerns sharing within the social circle. This result contributes to the empirical literature on peer effects in financial decisions and, in part, provides empirical support to theoretical models in which the decision to deposit and withdraw explicitly depends on the decision of others.

Appendix 1: Questionnaire

Our proxy for peers' financial sophistication is obtained from the following questions:

Q1 - *Which level of education do MOST of your acquaintances have*?

Q2 - *How much do you think is the AVERAGE total net income of those households*?

The information households' saving accounts is obtained from the following questions:

Q1 - *Who is the account owner of your (1st to 7th) saving or deposit account*?

Q2 - *With which bank or financial institution is your (1st to 7th) saving or deposit account registered*?

Q3 - *What was the balance on your (1st to 7th) saving or deposit account*?

Appendix 2: ATT magnitudes

Table A1.1: ATT from DiD

	2008	2009	2010	δ_1	δ_2
above	152.490,4	164.420,0	166.679,4	11.929,6	14.189
				(14.792,6)	(14.186,2)
below	39.552,1	39.762,6	42.194,5	210,5	2.642,4
				(1.701,9)	(1.697,1)
			(ATT)	11.720***	11.547***
				(930,1)	(802,5)

Note: The table reports estimated magnitudes of the Average Treatment Effects on the Treated (ATT). The top panel reports conditional group means, and δ_1, δ_2 represent the 2009 and 2010 within group changes with respect to the pre-reform year. The bottom panel reports the estimated ATT. The symbols *, ** and *** denote conventional significance levels.

Table A1.2: ATT from TD

$G=1$	2008	2009	2010	δ_1	δ_2
above	173.353,5	175.752,7	176.428,4	2.390,2	3.074,9
				(19.457,4)	(20.453,6)
below	39.711,0	38.896,8	42.823,0	-814,2	3.112,0
				(2.145,3)	(2.186,0)
			DiD_1	3.429**	-38,9
				(1.523,6)	(1.448,0)
$G=0$	2008	2009	2010	δ_1	δ_2
above	119.705,7	138.921,3	159.044,6	19.216	39.339
				(12.949,1)	(16.776,5)
below	39.232,6	41.763,0	42.157,0	2.530,4	2.924,4
				(2.805,7)	(2.671,2)
			DiD_2	16.685***	36.414***
				(1.549,3)	(1.568,2)
			(ATT)	-13.256***	-36.452,9***
				(650,3)	(471,0)

Note: The table reports estimated magnitudes of the Average Treatment Effects on the Treated (ATT). The top panel reports conditional group means, δ_1, δ_2 represent the 2009 and 2010 within group changes with respect to the pre-reform year and DiD_1, DiD_2 represent the 2009 and 2010 within group DiD. The bottom panel reports the estimated ATT. The symbols *, ** and *** denote conventional significance levels.

.

Chapter 2

Leverage Constraints, House Prices and Household Debt: Evidence from the Netherlands

2.1 Introduction

Understanding the origins and the causes of the Great Recession has been the main challenge for economists and policy makers in the last decade. Among the different explanations, a predominant view attributes the main cause to the "credit-driven demand channel", which caused a small problem in the US mortgage and housing market to trigger a worldwide financial meltdown.

Changes in models of banking origination first led to a strong increase in lending to households: in the U.S., mortgage debt almost doubled between 2000 and 2007 (Brown et al.

I am gratefully acknowledged to Mauro Mastrogiacomo for his supervision. I also thank Jakob De Haan, Gabriele Galati, Giacomo Pasini, Michele Pellizzari, Sweder Van Wijnbergen, Eduard Suari Andreu, Jurre Thiel and all participants at Vrije Universiteit Amsterdam, De Nederlandsche Bank, the NIBUD Institute, the 3rd Applied Microeconomics Workshop, the Dutch Economist Association (KVS), Ca' Foscari University of Venice, Netspar, the Dutch Economists Day (NED), University of Modena di Reggio Emilia, the 2019 Winter Meeting of the Econometric Society (ESWM) and the Ministry of Interior (BZK) for their useful comments. Eventually, I thank Frank van Hoenselaar, Dorinth van Dijk and the Research Department at De Nederlandsche Bank for providing me the data used in the empirical analysis. The views and opinions expressed in this article are those of the author and do not necessarily reflect the position of De Nederlandsche Bank.

2010), it increased across the whole income distribution (Adelino et al. 2016) and also among subprime borrowers (Mian and Sufi, 2009). The increase in lending then boosted household demand and led to an increase in house prices, which in turn had a feedback effect on household leverage through home equity based borrowing (Mian and Sufi, 2011) or expectations of higher house prices (Kaplan, Mittman and Violante, 2017). The initial increase in house prices rapidly transformed into a bubble, that lasted until the 2007-2008 crash, which led to undesired outcomes such as foreclosures (Mian, Sufi and Trebbi, 2015), defaults (Mayer, Pence and Sherlund, 2009) and consumption cuts (Mian, Rao and Sufi, 2013).

As a consequence, a valuable lesson from the Great Recession was that, to prevent future financial and economic crises, it is important to look at lending growth and household debt, and in particular to the housing finance component (Schularick and Taylor, 2014). Policy makers around the world thus took a stronger regulatory approach in this segment of the financial sector and, in particular, they have been increasingly relying on macroprudential policies to prevent excessive risk-taking in the financial system.

In this paper I investigate the effects of macroprudential policies, in the form of Loan-to-income (LTI) and Loan-to-Value (LTV) limits, on the debt amounts of newly-originated residential mortgages in the Netherlands, during a housing market boom.

Currently, the Netherlands is one of the most interesting countries in which to study the link between leverage constraints, house prices and households debt: according to the Oecd[1], the Netherlands is the country with the second highest level of mortgage debt in the world and, according to the International Monetary Fund[2], it has experienced one of the largest house price increases worldwide in the recent years. In fact, the Financial Stability Report[3] published by the Dutch Central Bank states that the link between the housing and the mortgage market is the main source of financial stability risks in the country. The introduction of leverage constraints was aimed precisely at containing households indebtedness, which has always been traditionally very high in the Netherlands[4].

[1]Oecd National account statistics, available at https://data.oecd.org/hha/household-debt.htm

[2]IMF Global Housing watch, available at https://www.imf.org/external/research/housing/

[3]Available at https://www.dnb.nl/en/news/news-and-archive/Persberichten2019/dnb384297.jsp

[4]In the pre-crisis period, it was pretty common to borrow mortgages with LTV ratios between 100% and 120%. In fact, most of the high household debt in the Netherlands is due to the mortgage debt component.

Estimating the impact of leverage constraints and house prices on household debt is challenging. In overheating housing markets, more and more expensive properties for sale would force liquidity constrained borrowers to borrow higher mortgage amounts, closer to the limit established by the macroprudential policy. Even without a change in the leverage limit, this would become more binding, but ultimately because of increasing house prices. However, the causality chain can even run in the opposite direction. For example, a shock to lending supply such as a relaxation of the leverage limits allow households to borrow larger amounts, and increasing household debt would fuel housing demand in the economy, ultimately boosting house prices[5].

In this paper, I use administrative loan level data merged to local house price indexes constructed from individual transactions data in the housing market, to evaluate how leverage constraints and house prices impact on household debt. The effect is estimated both parametrically and non-parametrically, to analyze the impact on the debt amount that households borrow, as well as on the distribution of the borrowed amounts relative to the corresponding income and property values, respectively.

In the parametric specification, I jointly estimate the effect of leverage constraints and house prices on household debt. I do so using unique exogenous, unanticipated and granular changes in LTI limits, and using local house price indexes to account for increases in house prices within municipality, as well as for the heterogeneity in house price growth rates across municipalities. To deal with the potential reverse causality between house prices and household debt, I rely on an instrumental variable approach in which local house prices in Dutch municipalities are instrumented using a proxy of the total housing supply elasticity in each municipality. My approach is close to the one proposed by Mian and Sufi (2011), that I complement using as instrument a proxy of the total housing supply elasticity, represented by the share of developable land and the share of unoccupied dwellings. Results show that LTI limits are binding on average, but the estimated change in household debt caused by a change in leverage constraints decreases along the income distribution. LTI limits are particularly binding for low income households, who often borrow at the LTI limit or qualify for the exceptions established by the regulation.

[5]Even in absence of any debt limit, house prices and household debt are likely to be jointly determined, for example by an omitted variable such as shock to expected income growth, as shown in Attanasio and Weber (1994), Muellbauer and Murphy (1997), Mian and Sufi (2011).

For this group, results indicate a full pass-through of changes in leverage limits, which translate into an equivalent change in household debt, as households use all their borrowing capacity. Results also provide evidence of the competing effect of (increasing) house prices: once house prices are instrumented, the estimate of this effect doubles in magnitude, and suggests that a 1% increase in house prices induces a 0.6% increases in household debt. Correspondingly, the effect of LTI limits slightly attenuates. Eventually, LTV constrained borrowers are also found to borrow higher LTI mortgages, but gradual LTV limit reductions did not reduce mortgage amounts, as the increase in house prices has been stronger than the limit reduction.

In the non-parametric specification, this paper contributes to the literature on bunching introduced by Chetty et al. (2011); Kleven and Waseem (2013) and recently grown in terms of applications to the mortgage market literature (De Fusco and Paciorek (2017), De Fusco, Johnson and Mondragon (2019), Best et al. (2018)). This approach exploits the presence of non linearities in agents' budget sets to retrieve an estimate of the response at that specific point of the budget set[6]. In particular, I first propose a stylized model of borrowing showing that the effect of changes in leverage constraints on household debt is proportional to the mass of agents bunching at the LTI or LTV limit. Then, I use a bunching approach to obtain a reduced-form estimate of the number of households effectively constrained by these limits. To identify the share of LTI constrained borrowers, I use the approach of Chetty et al. (2011) to obtain an estimate of the counterfactual LTI distribution that exploit the *comply or explain* feature of the LTI regulation, and I exploit the unique granularity of LTI limits to exclude possible confounding effects. To identify the share of LTV constrained borrowers instead, I extend the traditional bunching approach to evaluate the effect of progressive and gradual reductions in the LTV limit established by the regulation. The approach, similar in spirit to the one developed by De Fusco et al. (2019), exploits time variation (instead of cross-sectional variation) in the location of the kink in the intertemporal budget set, and accounts for confounding effects such as increasing house prices.

Results show that the progressive 6% reduction in the LTV limit established by the LTV

[6]The typical example is the labor supply response in quantiles of the income distribution where the marginal tax rate sharply increases (kink). The response (labor supply elasticity) was shown to be proportional to the bunching mass at the kink point. See Kleven (2016) for a review.

regulation has impacted on up to 45% of highly leveraged borrowers. Results confirm again that LTI limits are only binding for low income households, but further show the role of flexibility established by the comply or explain regulation, which is likely preventing extensive margin responses of LTI limit changes. The main policy implication I draw from this paper is that, in order to properly contain household debt, a macro-prudential regulation that combines LTI limits with a LTV limit is desirable. On the contrary, an LTV rule alone cannot properly contain household debt growth in periods of booms and busts in the housing market. While this is the case in the Netherlands, most countries in the world (especially among advanced economies, see Cerrutti, Claessen, Leven (2017)) rely either on LTI rules or, most commonly, on LTV limits as the only macro-prudential tool. The remainder of the paper is organized as follows: Section 2.2 provides institutional details about the macro-prudential regulation in the Netherlands, Section 2.3 introduces a stylized theoretical framework that motivates and guides the following empirical study. Sections 2.4 and 2.5 present the data, the descriptive and the empirical analysis. Section 2.6 concludes.

2.2 Institutional framework

2.2.1 Macro-Prudential Policy

A macro-prudential regulation is a policy framework that aims to limit risk intake in the financial system. The ECB is the main macro-prudential regulator in the Euro Area, but some macro-prudential policies remain under the control of national governments or national supervisory authorities (typically, national central banks). Among the different policies, LTV and LTI limits are the most common macro-prudential tools aimed at limiting excessive indebtedness of households and firms. These instruments represent the core of macro-prudential regulation also in the Netherlands.

The LTV limit establishes the maximum debt that can be lent to a borrower, relative to the collateral value of its house. The Dutch rule is straightforward: as of 2012 originated mortgage loan amounts must be at most equal to 106% of the appraised house value.

Then this limit has been reduced by 1% every year up to 2018, when the LTV ratio has been set permanently to 100%. Therefore, the LTV ratio displays time variation but no cross-sectional variation as it is the same for all borrowers.

The LTI limits are set by the Dutch government at the recommendation of the National Institute for Family Finance Innovation (NIBUD) as a debt-service-to-income constraint (DSTI). A DSTI constraint establishes the maximum debt service that a household can afford to pay on a monthly basis, as a percentage of its income. These account for all the necessary expenses that families incur into: they are based on budgeting computations that account, among others, for the level of consumer and energy prices, interest rates and taxation.Eventually, the recommended DSTI limits are converted into equivalent LTI limits, that establish the maximum amount that can be lent to a borrower, as a multiple of family income. In the empirical analysis, I will refer only to the resulting LTI limits, as these are the ones that banks apply at origination. The resulting LTI limits reflect the affordability of debt repayment, and thus depend on household income and on the interest rate, which determines the debt service.

Figure 2.1: Regulatory Loan-to-Income limits (Heatmap).

Bruto	Hypotheekrente				
jaarinkomen	3.75%	4.25%	4.75%	5.25%	5.75%
19500	3.0	2.9	2.8	2.7	2.6
20000	3.1	3.0	3.0	2.9	2.8
20500	3.3	3.2	3.1	3.0	2.9
21000	3.5	3.4	3.3	3.2	3.1
21500	3.6	3.5	3.4	3.2	3.1
22000	3.8	3.6	3.5	3.4	3.3
22500	3.9	3.8	3.8	3.7	3.6
23000	4.0	4.0	3.9	3.8	3.8
23500	4.1	4.1	4.0	3.9	3.9
24000	4.2	4.2	4.1	4.0	3.9
25000	4.4	4.3	4.2	4.2	4.1
26000	4.5	4.4	4.3	4.2	4.1
28000	4.6	4.5	4.4	4.3	4.2
55000	4.7	4.6	4.5	4.4	4.3
58000	4.8	4.7	4.6	4.5	4.4
61000	4.9	4.7	4.6	4.5	4.4
63000	4.9	4.8	4.7	4.6	4.5
65000	5.0	4.9	4.8	4.7	4.6
68000	5.1	5.0	4.9	4.8	4.6
70000	5.2	5.1	5.0	4.8	4.7
75000	5.3	5.2	5.0	4.9	4.8
77000	5.3	5.3	5.2	5.1	5.1
79000	5.4	5.3	5.3	5.2	5.1
85000	5.5	5.4	5.4	5.3	5.2
96000	5.6	5.5	5.4	5.4	5.3
110000	5.7	5.6	5.5	5.4	5.4

Note: The figure shows the heatmap of a table containing the recommended Loan-to-Income (LTI) limits. LTI limits depend on the gross annual household income (vertical axis) and on the interest rate charged on the mortgage (horizontal axis). Stricter limits are depicted in red, larger limits are depicted in green. The example is the actual 2014 LTI limit table. The LTI limits are expressed as multiples the borrower's gross annual household income.

Since the resulting LTI limits reflect the affordability of debt repayment, they depend on total household income and on the interest rate paid on the mortgage, which is part of the debt service. Figure 2.1 reports an example of the table containing the recommended LTI limits. The example refers to the rule in force in 2014. The figure shows the considerable cross-sectional variation that the limits display: depending on income and on the interest rate, they ranged from 2.6 to 5.7. An LTI limit of 4 indicates that the maximum loan amount that a household affords to repay is four times its gross annual income. Stricter LTI limits are assigned to lower-income and riskier (high interest rate) households. These LTI limits also display considerable time variation other than cross-sectional variation, as the recommendations are revised annually. Figure 2.2 shows the changes in the maximum borrowing capacity between the years 2015-2016 and 2016-2017, for the different income brackets.

Figure 2.2: Changes in maximum allowed indebtedness

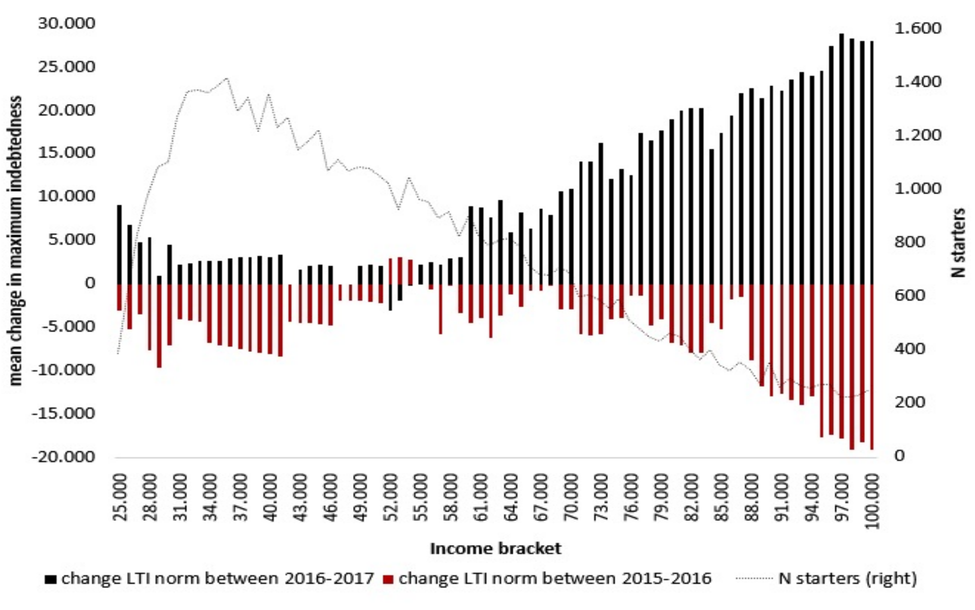

Note: The figure shows the average change in the maximum loan amounts in 2015-2016 (in red) and 2016-2017 (in black) respectively, for different household income categories. The change in each income category represents the average change among all interest rate categories in the same income category. Also the figure shows the empirical income distribution among starters household in the sample considered.

As an example, consider a household with an income of 50.000 euro. A 0.2 LTI limit change reflects a 10.000 euro change in the maximum loan amount. The Figure shows that between 2015 and 2017 changes in the limits have been both positive and negative, and that the resulting changes in the borrowing capacity have been sizable, especially for high income households.

Also, a key feature of the LTI rule in the Netherlands is that it is a *comply or explain* rule, i.e. there are established exceptions that allow for flexibility options for both lenders and borrowers. In particular, if borrowers can qualify into one of the established exceptions, banks are allowed to grant them a mortgage with an LTI higher than the established limit. Banks can exceed the LTI limit if one of the following conditions apply: (*i*) mortgage refinancing (*ii*) energy saving investments and (*iii*) bridge loans (for borrowers who move into a new house, before the old house is sold). In addition to these specific cases, lenders are generally allowed to exceed the LTI limit if the decision is substantially motivated and documented. A notable example is the case of an expected increase of capital or labor incomes[7]. If the documentation doesn't meet the necessary requirement, banks willing to exceed the LTI limits are simply not allowed to do so, and must reduce the originated loan up to the limit. On the contrary, banks are never allowed to originate a mortgage exceeding the original 106% LTV cap and the LTV limit in place every year is a strict rule for first-time buyers (starters).

To summarize, LTI and LTV rules represent complementary leverage constraints aimed at limiting excessive household indebtedness. The former is a rule that takes into account the affordability of debt repayment, allows for flexibility options and depends on household characteristics. The latter instead establish a common limit to everybody, and is much stricter for starters. Section 3 investigates the theoretical implications of the introduction of such rules.

2.2.2 Mortgage Market

The mortgage market in the Netherlands is dominated by the four largest banks that provide about the 80% of the total supply of mortgages. The remaining market share is controlled by small banks, pension funds and insurance companies that have recently entered the market (Kim and Mastrogiacomo, 2019 ; Thiel, 2020). Traditionally, loan amounts granted by Dutch banks have always been very generous. In the period before the crisis, it was fairly common to borrow mortgages having Loan-to-Values up to 120%.

[7]For further legal details, see Van't Hof (2017).

Part of this phenomenon is certainly attributable to institutional features that explicitly incentivize high nominal debts in households' balance sheets. Examples include a mortgage interest deduction up to 51%, and a public housing guarantee (NHG, *Nationale Hypotheek Garantie*) that insures banks and borrowers against negative home equity risk in case of borrowers' default[8]. For this reasons, a number of non-amortizing products (such as Interest-only mortgages) or products that allow for tax arbitrage (such as Savings, Insurance and Investment mortgages) gained popularity, and now represent the most common mortgage types (DNB Occasional Study 13-4). Despite the high debt levels of Dutch households, mortgage defaults have always been very low. This can be partly explained by the presence of a full recourse system that allows lenders to lay claims on borrower's assets, other than those pledged as a protection.

2.2.3 Real Estate Market

The real estate market in the Netherlands is characterized by supply shortage, due to zoning restrictions and the presence of one of the largest social rental housing sector in Europe. While the existing stock of houses is scarce, there are also obstacles to increasing housing supply: municipalities depend on revenues from land development, but often are forced to lower land price to make the construction of regulated rental housing profitable (DNB occasional study, 15-1). The price elasticity of housing supply in the Netherlands is in fact among the lowest in Europe (OECD, 2011). Recently, the combination of a tight housing supply with a persistently increasing housing demand driven by immigration flows and an increasing investor activity have caused the housing market to overheat: house prices have rapidly increased in the 2013-2019 period at a very fast pace compared to international standards (and reached a new historical peak) and price to income ratios have been constantly rising too (DNB, 2019).

[8]Borrowers who buy NHG and default on the mortgage due to job loss, divorce or partner's death or disability, are insured against the possibility that the proceeds of the house sale are not enough to cover the outstanding nominal debt at default. Note that the NHG premium consist of a one-off payment proportional to the debt amount at inception, but is independent on the value of collaterals.

2.3 Theoretical framework

This section introduces a stylized theoretical framework to investigate the effect of leverage constraints on household borrowing. The focus is thus on the households borrowing decisions at origination. The proposed model builds on Piazzesi and Schneider (2016) who provide a general framework that includes *housing* in a life-cycle model. In fact, the borrowing decision is simultaneous to the house purchase decision, so this has to be accounted for. Related analyses have been proposed in Brueckner (1994), Defusco and Paciorek (2017) and Stein (1995)[9].

This model has at least two key distinctive features: first, houses are assets that provide a non-tradable dividend, the housing service, which is a consumption good. Second, individuals derive utility from living in their house (the housing service) and the utility is increasing in house quality. Since the focus is on the debt origination used to finance the house purchase, I do not consider the existence of a rental market but I only look at starting homeowners. For the same reason, I do not explicitly consider houses as technologies that depreciate if essential maintenance is not performed, and thus I exclude home improvement decisions from the analysis. Eventually, since the borrowing decision of a starting homeowner is by definition a one-time decision, I do not explicitly treat time as in Brueckner (1994), Defusco and Paciorek (2017) and Stein (1995). In particular, I consider households living T periods and borrowing at a given point in time $t < T$.

In the economy there are N households. Household i has a discount factor β_i distributed over the support $[\underline{\beta}, \overline{\beta}]$, according to a cumulative density function $F(\beta)$. At time t each borrower i maximizes:

$$\tilde{U}_i = max_{c_t,h_t,m_t} \; U(g(c_t, s_t(h_t))) + \beta_i E_t[V(w_{t+1})] \tag{2.1}$$

$$s.t. \quad c_t + ph_t = w_t + m_t \quad ; \quad w_t = y_t$$

$$w_{t+1} = (w_t - c_t - p_t h_t)R + p_{t+1} h_t$$

[9] Brueckner (1994) studies the relation between the demand for mortgage and the interest rates on savings and mortgages. Defusco and Paciorek (2017) look at the interest rate elasticity of mortgage demand. Eventually, Stein (1995) investigates the role of down-payments in explaining fluctuations in the housing market when agents are second time buyers.

Households choose consumption, the housing quality and the mortgage size that maximize life-cycle utility. Life-cycle utility corresponds to current utility, which is derived via consumption c_t and the housing service $s_t(h_t)$, and the expected future value of wealth $E_t[V(w_{t+1})]$ which represents the utility derived by optimally behaving in the remaining $T-t$ periods, conditional on w_{t+1} (Brueckner (1994)). Current utility takes the following functional form:

$$U(g(c,s(h))) = log\big(c^{\alpha}\, s(h)^{1-\alpha}\big) \tag{2.2}$$

While $V(w_{t+1}) = log(w_{t+1})$. From the first budget constraint the household uses its current endowment, represented by labor income y_t and the mortgage loan amount m_t, to finance the consumption expenses and the house purchase. A house of quality h_t is worth $p_t h_t$ and provides a housing service $s_t(h_t)$. Since the utility of living in a house is increasing in house quality, higher quality houses proportionally deliver higher housing services. Therefore, I follow Piazzesi and Schneider (2016) and set $s_t(h_t) = h_t$. From the second constraint, the housing and mortgage choice affect the level of future wealth w_{t+1} which is the difference between the future asset value of the house $p_{t+1}h_t$ and the outstanding mortgage debt Rm_t. I assume the interest rate $R = 1 + r$ to be certain and agreed upon the mortgage contract, while I assume uncertainty over future house prices p_{t+1}. For simplicity, I don't consider the role of beliefs and I simply assume that that $p_{t+1} = p_t + \epsilon_t$ with $\epsilon_t \sim IID(0, \sigma^2)$.

Next, I introduce two macro-prudential limits: the first is a loan-to-income (LTI) limit that allows to borrow up to a given fraction of borrower's income, the second is a loan-to-value (LTV) limit that allows to borrow up to a given share of total house worth. Borrowers decisions are thus subject also to the following constraints:

$$m_t \leq \theta y_t \qquad\qquad m_t \leq p_t h_t(1-\delta) \tag{2.3}$$

Where θ and δ are the policy parameters that determine the level of the LTI and LTV limits, respectively. To solve the model, I follow Piazzesi and Schneider (2016) who propose a two stage solution approach to the problem. In the first stage households choose the house quality h_t that trades off housing expenditures and housing utility. In the second stage, conditional on the optimal house quality, households decide how much to consume and how much to borrow. In other words, conditional on the house chosen, households use the mortgage as a consumption smoothing device. In the analysis that follows, I focus on the second stage problem, taking the optimized house quality h_t as given to investigate the corresponding inter-temporal allocation decision represented by the financing choice[10]. This inter-temporal allocation depends on the time preference β_i that is heterogenous across households[11]. With no assumptions on the functional form of U, g and V and conditional on the optimal house quality h_t, the Euler equation of the unconstrained case, when neither the LTI nor the LTV constraint bind, is equal to:

$$U'[g(c_t, h_t)]g'(c_t, h_t) = \beta_i \, E_t[V'(w_{t+1})]R \tag{2.4}$$

It establishes the relation between current and future consumption in the optimal consumption path. Again, the household uses the mortgage loan not only to finance the house purchase, but also to reach the best possible resource allocation described by the Euler equation. Using the functional forms of U, g and V in eq. (2.2), and using $E_t(p_{t+1}) = p_t$, from the Euler equation I solve for m_t to obtain the household i unconstrained mortgage function:

$$m_i^u = \frac{p_t h_t(\alpha + R\beta_i) - R\beta_i \, y_t}{R(\alpha + \beta_i)} \tag{2.5}$$

Where m_i^u denotes the mortgage size of borrower i in the unconstrained case[12]. The desired level of debt depends positively on property valuation and negatively on household income. Next, consider the constrained case with both constraints active and let $\lambda[m_t - \theta y] = 0$ and $\eta[m_t - p_t h_t(1-\delta)] = 0$ be the Kuhn Tucker conditions for the LTI and LTV constrains. Under these leverage limits the constrained mortgage amount is:

$$m_i^c = min\big\{\theta y_t \, ; p_t h_t(1 - \delta)\big\} \tag{2.6}$$

[10]This approach, despite formally treated in Piazzesi and Schneider (2016), has been implicitly adopted also in Brueckner (1994). Instead, Stein (1995) takes the same approach in the opposite perspective and studies the house quality choice conditional on the available endowment.

[11]Please note that the discount factor is the only element of heterogeneity in the population, as households are assumed to share the same life-time utility function, as well as the same level of income. This assumption, despite being strong in general, matches the aim of our empirical analysis that looks at the households' financing choices conditional on the observed household characteristics.

[12]The subscript t has been dropped to ease the notation.

In words, the constrained loan size is equal to the leverage constraint that binds first. Rearranging, the LTI limit binds first if:

$$\theta < \frac{p_t h_t}{y_t}(1-\delta) \tag{2.7}$$

Where $p_t h_t / y_t$ is the price-to-income ratio associated to borrowers' house purchases. In this model, whether leverage constraints are binding or not ultimately depends on each individual unconstrained loan size m_i^u, and thus on each individual discount factor β_i which is the only element of heterogeneity in the population: more patient households require levels of future consumption that are higher than those of inpatient households who, given lifetime resources, take highly leveraged positions to increase current consumption vis-a-vis future consumption. I derive the breakpoint level of β by equating the unconstrained to the constrained mortgage functions, in formulas: $\beta^* : m_i^u = m_i^c$. Suppose the LTI binds first, then:

$$\frac{p_t h_t(\alpha + R\beta_i) - R\beta_i\, y_t}{R(\alpha + \beta_i)} = \theta y_t \tag{2.8}$$

Leading to:

$$\beta^* = \frac{\alpha(p_t h_t - \theta R\, y_t)}{R[y_t(1+\theta) - p_t h_t]} \tag{2.9}$$

The value of β^* identifies marginal borrowers, i.e. the households whose unconstrained and constrained mortgage size coincide, for given levels of the leverage constraints. It is important to stress that marginal borrowers are completely unaffected by the policy since they are always able to borrow their desired loan size, which simply equals the first-binding constraint.

Figure 2.4 provides a graphical representation of the effect of leverage constraints on household debt in the simplest two periods case with the terminal condition $w_{t+1} = c_{t+1}$. In absence of credit constraints, the intertemporal allocation of household i would locate in point $A = (c_t^*, c_{t+1}^*)$ that represents the optimal solution of the unconstrained case. If leverage constraints are introduced, the inter-temporal budget set features a discontinuity and the household gets constrained by the regulation and locates at point B. This allocation is a second-best corner solution as $U'_i < U''_i$. The same would not be true for household j that, being more patient than i (i.e. $\beta_j > \beta_i$), chooses a mortgage size lower than the limit that still allows reaching the unconstrained allocation in point C. The implication is that, in absence of leverage constraints, the mortgage distribution in the population is the same as the distribution of discount factors $f(m) = f(\beta)$ (dashed line), while in presence of leverage constraints the same distribution would feature a spike at the leverage limit. The size of the spike is proportional to the number of borrowers constrained by the regulation.

Figure 2.3: (Un)constrained optimal consumption and mortgage distribution.

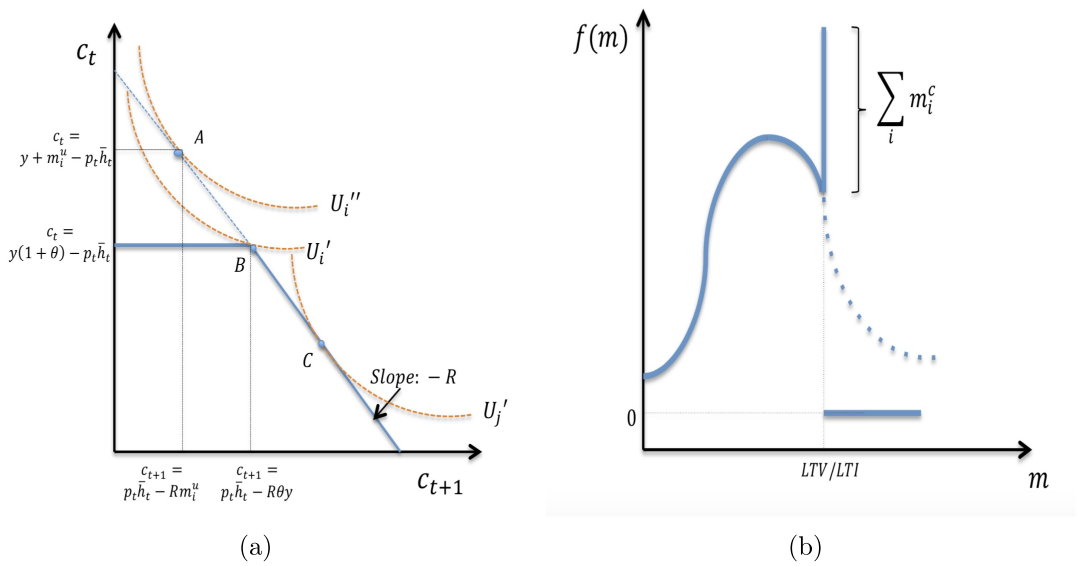

Note: The left figure shows the constrained (solid line) and the unconstrained (solid + dashed line) budget sets and the corresponding optimal solutions for current and future consumption. The right figure shows the constrained (solid line) and unconstrained (solid+dashed line) mortgage debt distributions.

2.3.1 Aggregation

In this section, I aggregate all households borrowing choices to determine the aggregate debt level in the population and its relation with the level of leverage constraints. According to the value of β^*, I can divide the population in two groups: the first group of unconstrained borrowers is made of patient households whose preferences imply taking low debt positions, these are all $i\ :\beta_i \in [\beta^*, \overline{\beta},]$. Conversely, the second group contains all constrained borrowers with discount factors $\beta_i \in [\underline{\beta}, \beta^*)$ that, being less patient, would tend to overindebt to maximize current utility in absence of leverage constraints. Again, the value of β is distributed according to a cumulative density function $F(\beta)$. In the former group, each individual debt level is different, as it depends on one's discount factor. In the latter group, everybody is constrained and takes the maximum allowed mortgage level. Let M^1 and M^2 be the corresponding aggregate group debt levels. Then, it follows that the average debt level in the population is the weighted sum of the debt levels in each group:

$$\begin{aligned} M(\theta) &= M^1 + M^2(\theta) \\ &= [F(\overline{\beta}) - F(\beta^*)]m_i^u + F(\beta^*)m_i^c(\theta) \end{aligned} \tag{2.10}$$

Where $F(\beta^*)$ is the share of constrained borrowers, which is increasing in the level of prices and decreasing in the level of income. From the last equation, the average debt level in the population explicitly depends on the policy parameter θ. The level of the leverage constraint affects not only the debt level of constrained households, but also the share of constrained borrowers in the population via the relation with $\beta^* = \beta^*(\theta)$. In fact:

$$F(\beta^*) = Pr(\beta_i \leq \beta^*) = F\left(\frac{\alpha(p_t h_t - \theta R\, y_t)}{R[y_t(1+\theta) - p_t h_t]}\right) \tag{2.11}$$

The stricter the leverage limit, the higher the share of constrained borrowers in the population. As a result, I can derive the change in the average debt level due to a change in the leverage constraint as:

$$\frac{\partial M(\theta)}{\partial\, \theta} = f(\beta^*)[m_i^c(\theta) - m_i^u] + F(\beta^*)\frac{\partial\, m_i^c(\theta)}{\partial\, \theta} \tag{2.12}$$

Where $f(\beta) = F'(\beta)$ is the probability density function. In words, the change in the policy parameter has two effects: on one side it changes the fraction of constrained and unconstrained borrowers in the population, on the other side it changes the aggregate debt level of constrained borrowers. Importantly, the change in the aggregate debt level caused by a change in the leverage limit is proportional to the bunching mass at the leverage limit, captured by $F(\beta^*)$: the higher it is, the higher the share of constrained borrowers and the larger the response to a leverage limit increase, as constrained borrowers' would use the policy change to increase their debt position. Note that eq. 2.12 holds for any possible density functions $F(\beta)$[13] and in the case the other constraint (LTV) binds first[14]. In a similar fashion, conditional on the same house bought[15], I can obtain the change in the aggregate debt level due to a change in house prices as:

$$\frac{\partial M(\theta)}{\partial\, p_t} = \begin{cases} \left[F(\overline{\beta}) - F(\beta^*)\right] \frac{\partial\, m_i^u}{\partial\, p_t} + f(\beta^*)\left[m_i^c - m_i^u\right] + F(\beta^*) \frac{\partial\, m_i^c}{\partial\, p_t} & if \quad \theta y_t > p_t h_t (1-\delta) \\ \\ \left[F(\overline{\beta}) - F(\beta^*)\right] \frac{\partial\, m_i^u}{\partial\, p_t} + f(\beta^*)\left[m_i^c - m_i^u\right] & if \quad \theta y_t \leq p_t h_t (1-\delta) \end{cases} \tag{2.13}$$

Interestingly, the effect of a house price change is different depending on which constraint binds first: since the borrowing capacity implied by the LTV limit $p_t h_t(1-\delta)$ is proportional to the level of house prices, an increase in prices does not preclude LTV constrained borrowers to increase the debt amount in line with house prices, as captured by the term $F(\beta^*)\ \partial m^c/\partial p$. In summary I obtain two main empirical implications from the model:

Empirical implication 1: Changes in house prices and changes in leverage limits jointly affect households' borrowing choice in terms of debt amount.

Following implication 1, an increase in house prices in a period of tightenings leverage constraints acts as an additional borrowing constraint. A combination of tighter limits and increases house price therefore crushes the distribution towards the limit. Conversely, in periods of constant borrowing constraints, an increase in prices pushes the level of debt towards the leverage limit, for reasons unrelated to the macro-prudential policy.

[13] In case $\beta \sim U[\underline{\beta}, \overline{\beta}]$ and the LTI limit binds first, eq. 2.12 has a closed form solution equal to:

$$\frac{\partial M(\theta)}{\partial\, \theta} = \frac{1}{\overline{\beta} - \underline{\beta}} \left(\theta y_t - \frac{p_t h_t(\alpha + R\beta_i) - R\beta_i\, y_t}{R(\alpha + \beta_i)} + \left(\beta^* - \underline{\beta}\ \right) y_t \right)$$

[14] In this case, the level and the change in debt will be denoted with $M(\delta)$ and $\partial M(\delta)/\partial\, \delta$, respectively.

[15] Note that a change in house prices should also affect the corresponding first-stage optimal house choice h^*. However, house qualities are often discrete and households might not be able to slightly "downsize" their housing choice.

Empirical implication 2: If leverage limits are binding, the distribution of household debt features a spike at the limit. The number of constrained borrowers and the elasticity of debt to a limit change are proportional to the spike.

Following implication 2, the aggregate effect on household debt of a policy aimed at changing the leverage limit will also be proportional to the bunching mass at the limit: in case of a limit increase, all borrowers take advantage of policy change to increase their debt position by an amount exactly equal to the increase in the limit. In case of a limit decrease, the bunching mass increase because of newly constrained borrowers, and because of intensive margin responses of already constrained borrowers.

2.4 Data and descriptive statistics

The main data source used in the empirical analysis is the Loan Level Data (LLD) collected by the Dutch National Bank. In order to comply with the 100% transparency policy of the ECB that allows banks to securitize their loans, banks must report all information required in the Residential Mortgage Backed Securities (RMBS) template of the ECB's European Data Warehouse. In the LLD, this information not only covers the pool of loans that banks plan to securitize, but refers to the entire mortgage portfolio of banks involved in securitization (see Mastrogiacomo and Van der Molen, 2015). Therefore, as of 2012Q4 banks report, on a quarterly basis, the information on all newly originated mortgage contracts (*recent production*) and, on a best effort basis, they also report information on mortgages granted before that date (*restrospective* part).

The information consists of borrower, property and loan characteristics for almost the 85% of the population of banks' issuing mortgage loans. The Dutch mortgage market is a very concentrated market in which the main three banks (ABN Amro, Rabobank and ING) control the largest share of the market. The activity of these three largest banks is well reported in the LLD.

I merge the LLD with three other data sources. The first and main data source comes from the Dutch Association of Real Estate brokers (NVM). This data contains house price indexes at the local level. In particular, these house price indexes are constructed using individual transaction data in the housing market, they are at quarterly frequency and at the two-digits postcode level (Van Dijk, 2019). In the Netherlands, the two-digits postcode unit approximately represents the municipality: the four biggest cities have a unique two-digit postcode, while in countryside areas the same two-digits postcode can be shared by two or more towns of the same province. Thanks to the granularity and frequency of these local house price indexes, I can account for quarterly differences in house price growth rates within and across municipality. Eventually, the last two data merged consist of municipality level information from Statistics Netherlands and the NIBUD Tables, a large file containing all the LTI limit recommendations set by the NIBUD Institute in the period 2012-2018.

Table 2.1 reports descriptive statistics on the most important property, borrower and loan characteristics in the LLD. The reported information only refers to first-time buyers (*starters*)[16]. The table shows that sustained increase in household debt over the sample years, in parallel with the increase in property valuations. Part of the increase in property valuations certainly reflects the sustained increase in the house prices evident in the NVM data. Table 2.2 shows that the increase in prices was generally high but varied substantially across regions: while the national house price index increased by almost 15% in a four year period, house prices have been increasing by 8% in the province of Drenthe and by nearly 50% in Amsterdam. The following analysis aims at explaining the reasons behind the increase in household debt, with a particular focus on the role of the macro-prudential regulation and increasing house prices.

[16] We focus on First-time buyers because this group is the one that leverage constraints explicitly target. Due to a reporting issue, the data does not allow observing renegotiating and starting borrowers in the data. To identify starters, I exploit a regulation, announced together with the LTV limit reduction, that establishes that as of 2013, the only mortgage types eligible for mortgage interest deduction (*hypotheekrenteaftrek*) are annuity and linear mortgages. Due to the generous tax deduction, other mortgage types disappeared from the market. This rule applies to newly originated mortgages, while borrowers who took their loan before 2013 and holding other types of mortgages are still eligible for the mortgage interest deduction. Therefore, we identify starters as borrowers whose mortgage has been originated and firstly reported after 2013, and whose mortgage type is either linear or annuity.

Table 2.1: Descriptive statistics (LLD)

		2014	2015	2016	2017	2018
Mortgage Debt						
	Mean	177.312,7	189.895,6	206.336.2	231.812,1	254.126.6
	Med	161.200,0	171.700,0	182.500,0	199.475,0	224.000,0
	N	52.251	56.575	62.530	67.807	57.426
Property Valuation						
	Mean	218.305,8	234.425,3	255.414,7	291.610,9	293.955,3
	Med	185.000,0	198.000,0	215.000,0	235.000,0	250.000,0
	N	52.005	55.760	62.095	67.665	57.426
Interest Rate						
	Mean	0,035	0,028	0,024	0,022	0,023
	Med	0,036	0,028	0,023	0,022	0,022
	N	52.005	55.760	62.095	67.665	57.426
Maturity						
	Mean	29,3	29,3	29,4	29,3	29,3
	Med	30	30	30	30	30
	N	52.251	56.575	62.530	67.807	57.426
Household Income						
	Mean	53.779,6	56.771.2	61.884.5	64.428.3	60.000,2
	Med	44.443,3	46.617,6	50.423,0	52.876,8	52.976,2
	N	52.251	56.575	62.530	67.807	57.426
Loan to Income						
	Limit(avg)	4.81	4.70	4.69	4.81	4.76
	Mean	3.63	3.67	3.65	3.87	
	Med	3.8	3.9	3.9	4.1	4.3
	N	52.132	56.739	62.509	67.406	57.426
Loan to Value						
	Limit	104.0	103.0	102.0	101.0	100
	Mean	85.9	86.7	86.2	84.2	84.3
	Med	98.2	98.7	97.5	95.8	98.9
	N	51.808	55.509	61.799	67.406	57.426

Note: Descriptive statistics at loan and borrower level. The top panel reports mean and median loan characteristics at origination: the debt amount at origination, the collateral value (property valuation), the interest rate and the maturity. The bottom panel reports borrower characteristics such as the mean and median LTI and LTV ratios, as well as the average LTI limit and LTV limit. The table eventually reports the number of observations for each variable.

Table 2.2: Descriptive statistics (NVM)

	2014	2015	2016	2017	2018		2014	2015	2016	2017	2018
National	100.9	103.7	109.0	117.2	125.6	Amsterdam	109.8	120.5	136.8	156.0	170.5
Drenthe	99.4	100.9	103.0	108.5	115.7	North Brabant	102.5	104.5	108.7	114.0	116.3
Flevoland	103.5	104.9	109.9	119.1	130.4	North Holland	103.3	108.7	117.7	130.2	139.5
Friesland	96.2	98.1	101.9	108.0	115.0	Overijssel	98.1	100.3	104.2	110.5	119.4
Gelderland	98.0	99.8	103.5	110.3	116.6	South Holland	102.5	105.5	110.9	120.2	119.3
Groningen	99.9	102.6	107.4	113.5	119.8	Utrecht	101.6	105.5	112.2	122.4	129.6
Limburg	102.4	104.4	108.6	113.9	120.5	Zeeland	105.9	106.4	109.0	112.0	118.6

Note: The table reports the value of the house price indexes at the provincial level for the period 2014-2017. In the top of the table, the national house price index and the local house price index in the municipality of Amsterdam are reported. The base year is the value of the national house price index in 2013.

2.5 Empirical analysis

The empirical analysis is divided in two parts. The first part tests Empirical Implication 1 via a joint estimation of the effects of house prices and LTI and LTV constraints on the amount of household debt at mortgage origination. The second part tests Empirical Implication 2 via a non-parametric estimation of the effects of LTI and LTV constraints on the corresponding distribution of the LTI and LTV ratios, respectively.

2.5.1 The effect of house prices, LTI and LTV limits

The first empirical implication drawn from the theoretical framework is that both leverage constraints and house prices affect household debt at mortgage origination. In particular, house price changes and changes in the leverage limits have a competing effect. Increasing house prices in a period of tightenings leverage constraints act as an additional borrowing constraint: when households' borrowing capacity decreases, higher house prices induce liquidity-constrained households to borrow more to purchase more expensive properties on sale. Conversely, in periods of constant debt limits, an increase in house prices pushes the level of debt towards the leverage limit, for reasons unrelated to the macro-prudential policy. This section tests this empirical implication, and aims at jointly estimating the effect of leverage limits and house prices on the level of household debt at origination.

To elicit the causal effect of the LTI regulation on household debt, I exploit the granularity of LTI limits, and their cross-sectional and time variation showed in Figures 2.1 and 2.2. Importantly, this variation is also exogenous, as the LTI recommendations are made by an independent institute on the basis of budgeting computations that account for changes in macroeconomic conditions such as changes in consumer prices, taxations and interest rates. Lastly, these variations can neither be anticipated nor foreseen: first, they cannot be anticipated as the LTI limits become effective in January, but become public only two months earlier, thus not allowing the time to purposely anticipate the underwriting of a mortgage, in case stricter limits are announced. Second, despite some of the changes in the macroeconomic conditions are easily predictable (e.g. the decrease in the interest rates due to the monetary policy stance), the resulting LTI limits are hard to predict because of the changing classification made by the NIBUD Institute during the sample period.

Figure 2.4: Exogenous changes in the income classification

Note: The figure shows the change in the number and the size of the income brackets undertaken by the NIBUD Institute for the LTI limits classification. The numbers and the dashes in black denote the different income brackets. The numbers in red denote the corresponding LTI limits, for the most frequent interest-rate category (3.0-3.5%).

Figure 2.4 clarifies this point. The NIBUD institute decided to gradually switch from broad income classifications (a total of four income brackets in a 30.000 euro income interval in 2014) to a very granular classification in 2017 (having one income bracket every 1.000 euro). As an example, two borrowers earning 30.000 and 55.000 euro would have been able to borrow an amount equal to 4.7 times their own income in 2014. As a result of the change in classification, borrowers with the very same income in 2016 would be subject to a 4.4 and 4.8 LTI limit respectively. Therefore, borrowers who share the same LTI category in a year can be subject to both tightening and loosening the next year, depending on the new income classification and their relative income position.
Given these premises, I estimate the following specification:

$$log(Mortgage\,amount)_{i,t} = \beta_1 \; LTI^{max}_{i,t} + \beta_2 P_{m,t} + \mathbf{X}_{i,t}\delta + \epsilon_{i,t} \tag{2.14}$$

Where the dependent variable is the amount of mortgage debt taken out by borrower i in year t, expressed in logs. The main coefficient of interest is β_1 that captures the effect of the LTI limit on the level of debt, while $\mathbf{X}_{i,t}$ is a $(n \times k)$ matrix of borrower, property and loan characteristics. These include the household gross annual income, the house type, the loan type, the mortgage interest rate and maturity, the employment status, an indicator for whether the borrower lives in a big city and an indicator for whether the borrower is covered by the National Housing Guarantee. $P_{m,t}$ is house price index in municipality m in year t, and represents the key conditioning variable to account for the competing effect of house prices on household debt. Eventually, eq. 2.14 is augmented with bank, time, region and region-time fixed effects and estimated on a pool of repeated cross-sections[17].
To estimate eq. 2.14, I propose an identification strategy based on instrumental variables (IV) to deal with the potential simultaneity issue between household debt and house price growth. In fact, the causality between house prices and household debt may even run in the opposite direction than what eq. 2.14 shows: for instance, a positive shift in credit supply may induce households to take on more debt, and this would in turn boost housing demand and house prices in equilibrium.

[17]Since we look at the mortgage origination of starting homeowners we never observe the same borrower in different periods, as this borrowing and house purchase decision is by definition taken once in a life-time.

I instrument the local house price index using two proxies of the supply and the elasticity of supply in the Dutch local housing markets in 2013. The first instrument is the share of developed land[18] introduced by Saiz (2010) studied in Hilber and Vermeulen (2016), and firstly used as instrument by Mian and Sufi (2009) for the United States. The intuition is the following: for a given shock to housing demand, the equilibrium price in the housing market should clear at higher levels in municipalities characterized by mostly urbanized and developed areas. The higher the share of developed land, the lower the possibilities for urban expansions. The second instrument is the share of unoccupied dwellings, which instead proxies the existing stock of houses: for a given shock to housing demand, house prices should clear at lower levels in cities characterized by excess supply of existing dwellings. Taken together, the two instruments proxy the elasticity of housing supply in each municipality.

Figure 2.5: Instrumental Variables

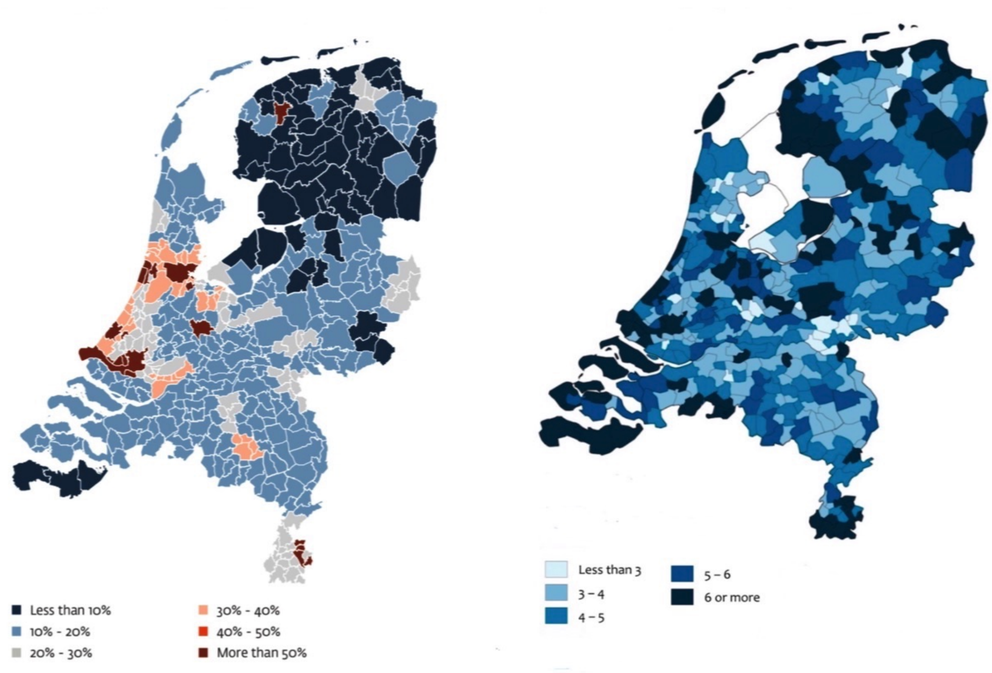

Note: The Figure shows the geographical variation in the share of developed land (left) and in the share of unoccupied dwellings (right). Source: CBS, DNB.

[18] The share of developed land is taken from the land use classification provided in the Land Cover Map. It is defined as the size of the developed land over the total developable land. Water, despite being developable in the long run, is excluded from the total developable land which includes mostly fields, grass and woods.

Figure 2.5 shows the geographical variation of these two instruments: it shows that in the four big cities (where house prices are growing at the highest pace, see Table 2) the share of developed land is above 50% and possibilities of urban developments are therefore very limited. On the contrary, the share of unoccupied dwellings tend to be larger in the north and in the south-western part of the Netherlands, which has more villages, smaller towns, and rural locations than the more-densely populated Randstad area. I estimate specification 2.14 using Optimal GMM, which is more efficient than 2SLS in the over-identified case. Results are reported in Table 2.3, that include also the OLS estimates of the same specification.

Table 2.3: The Effect of Leverage Constraints and House Prices on Household Debt

	Dependent variable: loan amount					
	OLS	IV-GMM	OLS	IV-GMM	OLS	IV-GMM
LTI limit	0.348***	0.254***	0.396***	0.335***	0.4403***	0.374***
	(0.004)	(0.013)	(0.005)	(0.010)	(0.005)	(0.005)
LTI limit × income			-0.0026***	-0.0039***	-0.0023***	-0.0037***
			(0.0001)	(0.001)	(0.0002)	(0.0003)
LTI limit × LTV constr.					0.028**	0.035***
					(0.0033)	(0.001)
Log Local house price index	0.287***	0.658***	0.284***	0.655***	0.0012***	0.609***
	(0.012)	(0.083)	(0.013)	(0.082)	(0.013)	(0.081)
controls	✓	✓	✓	✓	✓	✓
bank FE	✓	✓	✓	✓	✓	✓
time FE	✓	✓	✓	✓	✓	✓
region FE	✓	✓	✓	✓	✓	✓
Hansen's J (overid test)	-	0.0012	-	0.0022	-	0.0007
N observations (Nt)	263.594	259.533	263.594	259.533	263.594	259.533

Note: The dependent variable is the log of the borrowed loan amount. The estimates of columns (1), (3) and (5) are Pooled OLS, while the estimates in columns (2), (4) and (6) are IV GMM estimates (Optimal GMM). Robust standard errors in brackets. The set of covariates include loan characteristics (NHG, interest rate, maturity), borrower characteristics (income and age), a set of employment status dummies and a set of house type dummies and an indicator for urban areas. The symbols *, **, and *** denote conventional statistical significance levels.

Regarding the LTI rule, results from Table 2.3 show that the level of LTI limits explains the amount of debt at origination and that, on average, LTI limits are binding. The coefficient magnitudes from the first specification suggest that, on average, a 0.1 LTI limit change is associated to a 2.5%-3.5% increase in household debt. The second specification shows that this effect is highly heterogeneous along the income distribution. Figure 2.6 reports the estimated marginal effect of a 0.2 LTI limit change and shows that the LTI limit is particularly constraining for low-income households: for a EUR 30.000 income household, a 0.2 LTI limit increase is associated to an increase in the borrowing by about EUR 6.000, which corresponds to the increase in the borrowing capacity induced by the policy change. This result indicates a complete pass-through of a change in borrowing capacity into a change in household debt for very low income households. On the contrary, changes in debt are totally decoupled from changes in the LTI limit at very high incomes, while for the mid income categories an increase in the LTI limit causes an increase in debt by a positive fraction. Eventually, the third specification shows that the two limits tend to bind jointly, as LTV constrained borrow 2.8%-3.5% higher LTI mortgages.

Figure 2.6: Response to a 0.2 LTI limit change

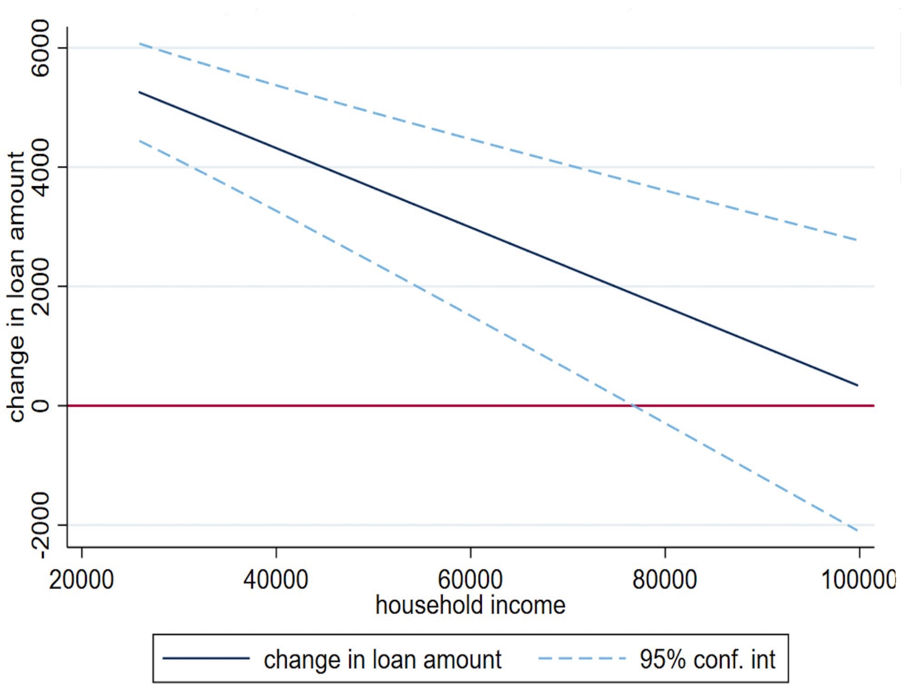

Note: The figure shows the marginal effect of a 0.2 LTI limit change, as a function of the annual household income. The marginal effect is averaged over the sample.

Regarding house prices, results provide evidence that house prices acted as an additional binding factor in households' borrowing decisions. The coefficient of the IV specification suggests that a one percent increase in house prices is associated to an increase in household debt by around 0,6%. Also, results show that once house prices are instrumented, the size of the IV coefficient is two times higher than in the OLS specification. Correspondingly the effect of the LTI limit reduces by about one third, relatively to OLS. To compare the relative strength of the effects of house prices and LTI limits, I compute the standardized coefficients of the OLS and IV specification. The first suggest that a standard deviation increase house prices and in the LTI limit is associated to an increase in household debt by 8% and 29%, respectively. In the IV specification instead, the two effects are much more similar in magnitude, and equal to 18.5% and 20% respectively. The comparison between the OLS and IV specifications highlights the intuition that, once the endogeneity issues have been addressed, the role of the increase in house prices emerges in pushing debt levels towards regulatory limits.

2.5.2 Bunching at the LTI limit

The second implication drawn from the model is that, if LTI limits are binding, we should observe a spike at the limit in the corresponding distribution of LTIs, and the size of this spike is informative on how binding LTI limits are.

Contrarily to the previous section, here I don't look at the effect of the LTI regulation on the amount of debt at origination, but I look at the effect on the LTI distribution. The aim is to understand how many borrowers are constrained by the LTI regulation, and to shed light on the role of the flexibility options that a *comply or explain* regulation allow to borrowers and banks.

With a strict LTI regulation that does not allow any flexibility, all constrained borrowers switch from the unconstrained choice to the constrained choice in which they borrow at the LTI limit. As depicted in Figure 2.3, this creates a spike in the mortgage distribution. With a comply or explain regulation, a third group of borrowers is likely to emerge: similarly to constrained borrowers, this new group of borrowers demands a loan amount that exceeds the limit but, unlike constrained borrowers, they are granted it because they qualify for one of the exceptions established by the regulation. Therefore, borrowers in this group are *de facto* unconstrained.

I evaluate how binding LTI limits are by estimating the size of the spike in the LTI distribution[19]. The challenge in estimating this is that it is possible to observe the mortgage distribution either when LTI limits are in force or when they are not, but both distributions are never observed jointly. To overcome this, I rely on an estimate of the counterfactual distribution. This counterfactual distribution is estimated under the assumption that most of the effect of the LTI regulation on the LTI distribution is local, at the LTI limit. In particular, this counterfactual distribution matches the observed distribution away from the LTI limit area, but not at the limit, where the corresponding density is obtained as a smooth interpolation between the densities at the right and the left of the limit.

In such a way, I obtain an estimate of the distribution that I would observe in absence of the leverage limit, and that accounts for the presence of marginal borrowers. I do this by following the bunching approach first introduced by Chetty et al. (2012) and Saez (2012). I follow the approach of Kleven (2016) and estimate the following specification:

$$n_j = \sum_{i=1}^{p} \beta_i (z_j)^i + \sum_{s=-k}^{+k} \delta_s \mathbf{1}\left[z_j = z_{c+s}\right] + \epsilon_j \tag{2.15}$$

Where n_j is a count variable for the number of loans in each bin j. The running variable z_j is the bin count obtained by discretizing the distribution is J equally-spaced bins. In formulas: $z_j = [1, ..., z_c, ..., J]$ where z_c is the LTI limit bin.

As a result, the first part of eq. 2.15 is a p-degree polynomial fit of the distribution of LTI. The second term of the equation contains a set of dummies that take value one for all bins in a window of size $2k$ around the LTI limit bin z_c. This term captures the local feature of the LTI distribution due to the presence of LTI-constrained borrowers. Assuming smoothness of the true counterfactual distribution[20], the estimated counterfactual distribution

[19] As clear from the theoretical model, to answer this question it is not enough to look at the share of households borrowing at the LTI limit. This is because marginal borrowers, who also borrow at the limit, are unconstrained by the regulation.

[20] Equivalently, the assumption is that without the LTI regulation, the LTI distribution would have been smooth in the LTI limit region, i.e. it would not display any spikes or discontinuities. Equivalently, the assumption states that the spike at the LTI limit in the observed distribution is solely attributable to the LTI regulation.

is obtained as the predicted value of eq. 2.15 omitting the contribution of the dummies, that is: $\widehat{n_j} = \sum_{i=1}^{p} \beta_i (z_j)^i$. This estimate provides a distributional fit based on the whole shape of the empirical distribution, but the local feature at the limit[21].

Let B be the excess number of loans at the LTI limit, then:

$$\begin{aligned} \widehat{B} &= n_j - \widehat{n_j} \\ &= \sum_{s=-k}^{+k} \delta_s \mathbf{1}\big[z_j = z_{c+s}\big] \end{aligned} \tag{2.16}$$

The estimates $\widehat{B}$ is an absolute measure of bunching, and is measured in number of loans. I also obtain the relative measure of bunching $\widehat{b}$ by scaling the previous statistic by the average density in the LTI constraint area:

$$\widehat{b} = \frac{\widehat{B}}{\sum_{s=-k}^{k} \widehat{n}_{c+s}/(2k+1)} \tag{2.17}$$

The statistics $\widehat{B}$ and $\widehat{b}$ represent reduced form non-parametric estimates of the number and the share of constrained borrowers, respectively.

Due to the presence of multiple LTI limits, we group all borrowers sharing the same LTI constraint and we run separate bunching estimates at all LTI limits from 4.4 to 5.3[22]. Also, I choose the free parameters J, p and k of eq. 2.15 in such a way that the resulting estimate is as much conservative as possible[23].

I report the bunching estimates for the LTI distributions in Table 4. I compute the standard errors using a parametric bootstrap procedure in which I draw with replacement from the raw LTI distribution and I re-compute the parameters at each bootstrap replication.

[21] Please note that is thanks to the comply or explain nature of the LTI rule that we can obtain a smooth counterfactual density in the LTI limit area. In fact, it ensures the presence of enough mass at the left (comply) and the right (explain) of the LTI limit, so that the resulting interpolation is likely to be smooth.

[22] Since this approach requires large data, we choose this interval in such a way that all estimates have at least 10.000 observations. This selection excludes borrowers in the tails of the income distribution, but includes almost 90% of all borrowers in our sample. For the same reason, I pool all borrowers subject to the same LTI limit over the whole sample period.

[23] The LTI distribution is discretized in bins of width 0.05 and the analysis area (J) includes a window of 60 bins. For the bunching area (k) I take a window of 1 bin around the LTI limit. Eventually, in line with the literature, the estimate of the counterfactual distribution is based on a 7th-degree polynomial (p).

Table 2.4: Bunching at the LTI limit

Limit = 4.4	est.	95% conf. int.	**Limit = 4.9**	est.	95% conf. int.
$\widehat{B}$	245.6***	[+149.2 ; +367.5]	$\widehat{B}$	1.7	[-56.0 ; +52.6]
$\hat{b}$	0.703***	[+0.416 ; +1.101]	$\hat{b}$	0.026	[-0.779 ; +0.819]
N	10705		N	15.887	
Limit = 4.5	est.	95% conf. int.	**Limit = 5.0**	est.	95% conf. int.
$\widehat{B}$	1484.9***	[+1297.5 ; +1699.5]	$\widehat{B}$	18.3	[-33.1 ; +52.5]
$\hat{b}$	1.166***	[+1.007 ; +1.340]	$\hat{b}$	0.470	[-0.733 ; 1.485]
N	39.552		N	12.364	
Limit = 4.6	est.	95% conf. int.	**Limit = 5.1**	est.	95% conf. int.
$\widehat{B}$	2727.5***	[+2489.4 ; +3005.9]	$\widehat{B}$	16.9	[-44.2 ; +69.6]
$\hat{b}$	1.634***	[+1.477 ; +1.829]	$\hat{b}$	0.269	[-0.618 ; +1.198]
N	55.984		N	10.131	
Limit= 4.7	est.	95% conf. int.	**Limit = 5.2**	est.	95% conf. int.
$\widehat{B}$	1036.5***	[+787.1 ; +1272.4]	$\widehat{B}$	10.2	[-45.6 ; +58.9]
$\hat{b}$	0.726***	[+0.540 ; +0.906]	$\hat{b}$	0.136	[-0.547 ; +0.827]
N	60.117		N	17.620	
Limit = 4.8	est.	95% conf. int.	**Limit = 5.3**	est.	95% conf. int.
$\widehat{B}$	-177.2	[-260.8 ; -78.1]	$\widehat{B}$	56.8	[-13.3 ; +119.6]
$\hat{b}$	-0.846	[-0.906; -0.540]	$\hat{b}$	0.680	[-0.147 ; +1.529]
N	15.887		N	15.794	

Note: The table reports pooled bunching estimates across all LTI limit categories, for the whole sample period. It reports the absolute ($\widehat{B}$) and relative ($\hat{b}$) bunching mass in each LTI limit class, as well as the corresponding 95% bootstrapped confidence intervals. Also, the table reports the total number of borrowers (N) in each LTI limit class. The symbols *, **, *** denote conventional statistical significance levels.

Table 4 shows significant bunching estimates in the mortgage distributions of all borrowers subject to LTI limits below 4.8. The LTI class in which most borrowers are LTI-constrained is the one with an LTI limit of 4.6. In this category, 2700 out of a total of 55.978 borrowers bunch at the limit, and the estimated bunching mass is about 60% higher than the average counterfactual density in the LTI limit region. Similarly, borrowers subject to nearby LTI limits seem to be strongly affected by the regulation: the bunching estimates are strongly significant for borrowers subject to LTI limits in the interval [4.5 , 4.7]. These three categories represent more than 40% of all borrowers in our sample and apply to below-average income borrowers earning less than 60.000€, as clear from Figure 2.1. Figure 2.7 provides a graphical representation of the results in Table 2.4. The area within the two dashed red lines represents the analysis area, and results from the choice of parameter J. The area within the two green lines represents the LTI limit area, which corresponds to the window of size $2k$ around the LTI limit bin. The thick black line is instead the estimated counterfactual distribution.

For low LTI limits, we observe that the estimated distribution is way below the actual distribution in the limit region, and the corresponding bunching mass is very high. At high LTI limits instead, the estimated distribution perfectly fits the empirical distribution in the LTI limit region, and the corresponding bunching mass is thus almost invisible.

That means that richer households, who are subject to higher LTI limits, are relatively unaffected by the policy: they tend to borrow less than they could, and the share of explainers is very low. On the contrary, we see that the share of explainers is very high at low LTI limits. This suggests that low income households, who also lack assets to make large downpayments, participate to the credit market mostly conditional on qualifying for one of the exceptions established. This also suggests possible extensive margin responses to changes in LTI limits[24]. Eventually, the high heterogeneity in the LTI limits, together with evidence of bunching at different LTI levels, let me reasonably exclude the possibility that bunching at a given LTI are due to other reasons, unrelated with the regulation.

In summary, the number of constrained borrowers is high only among low-income households. Still, the flexibility offered by the regulation has a role in allowing many marginal borrowers to exceed the LTI limit when constrained.

[24]Lacking data on loan applications, I cannot investigate this aspect.

Figure 2.7(a): Bunching at the LTI distribution

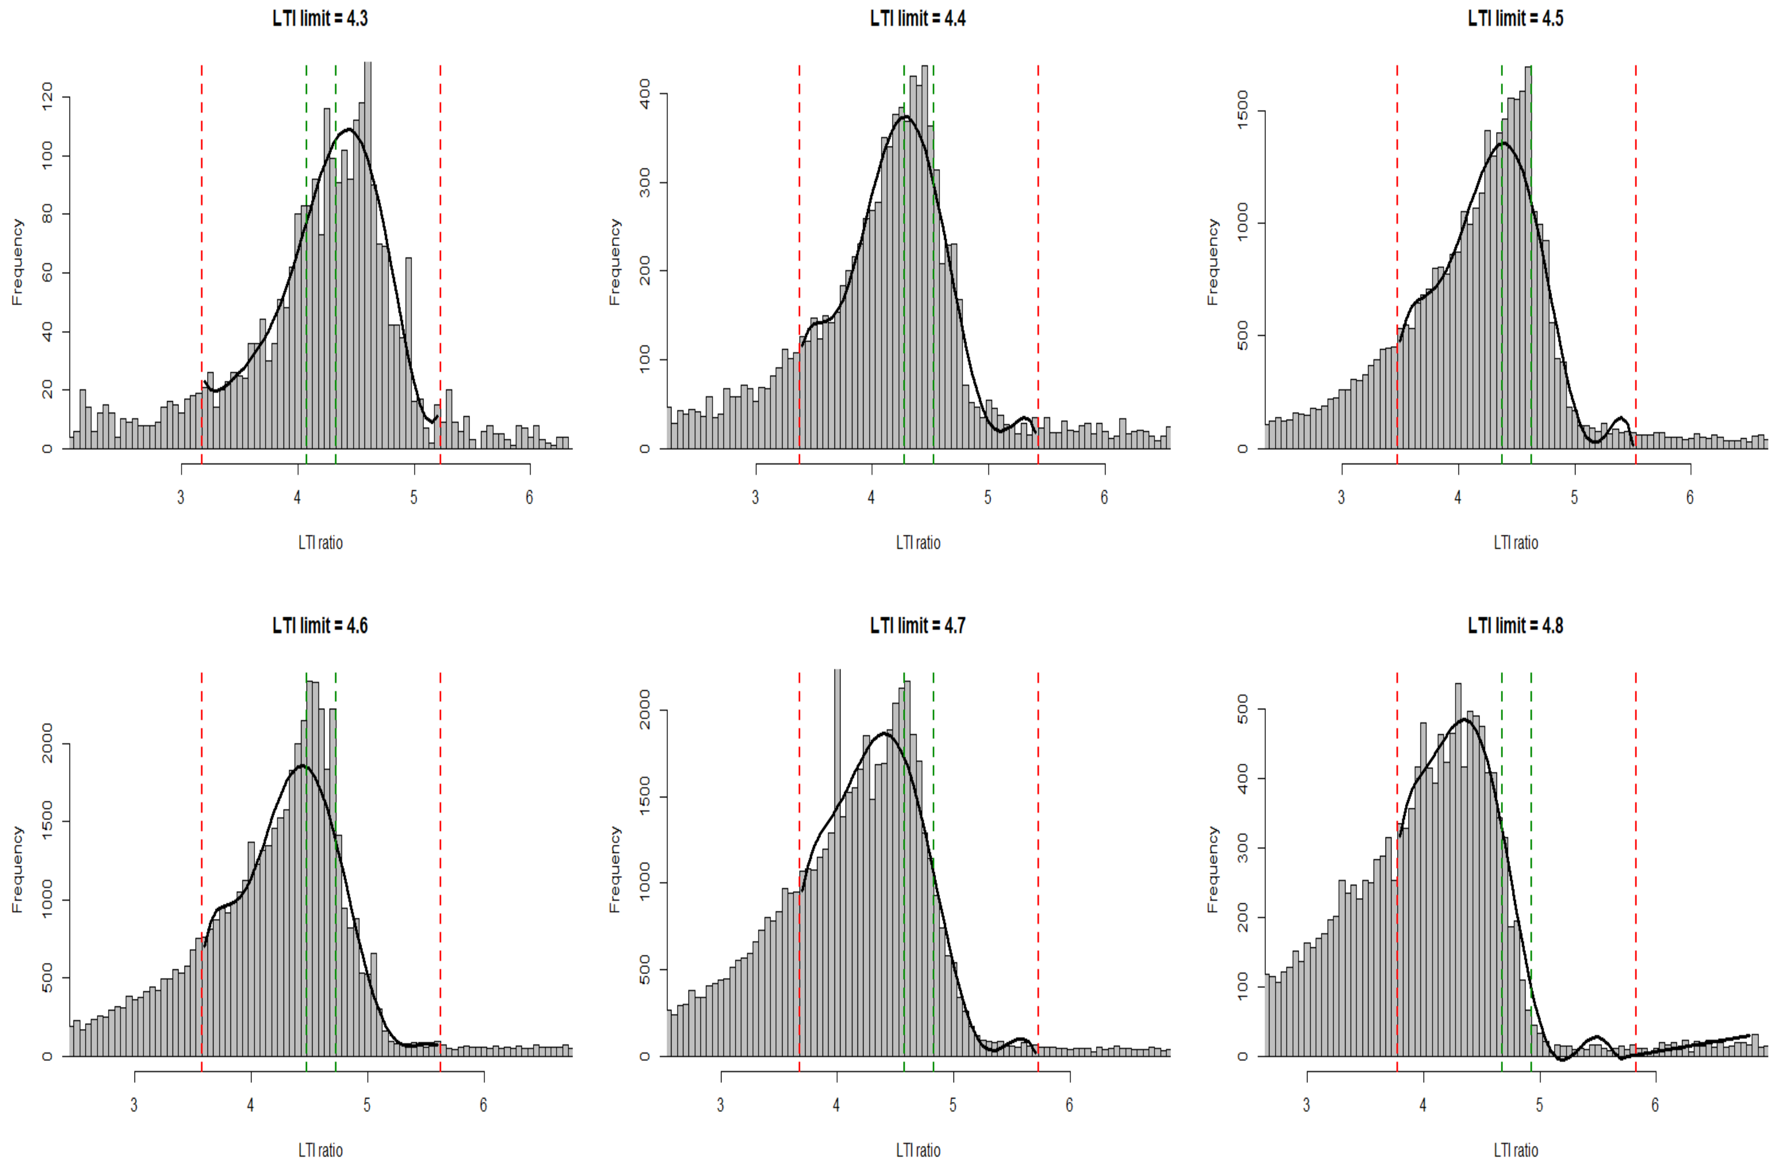

Figure 2.7 (b): Bunching at the LTI distribution

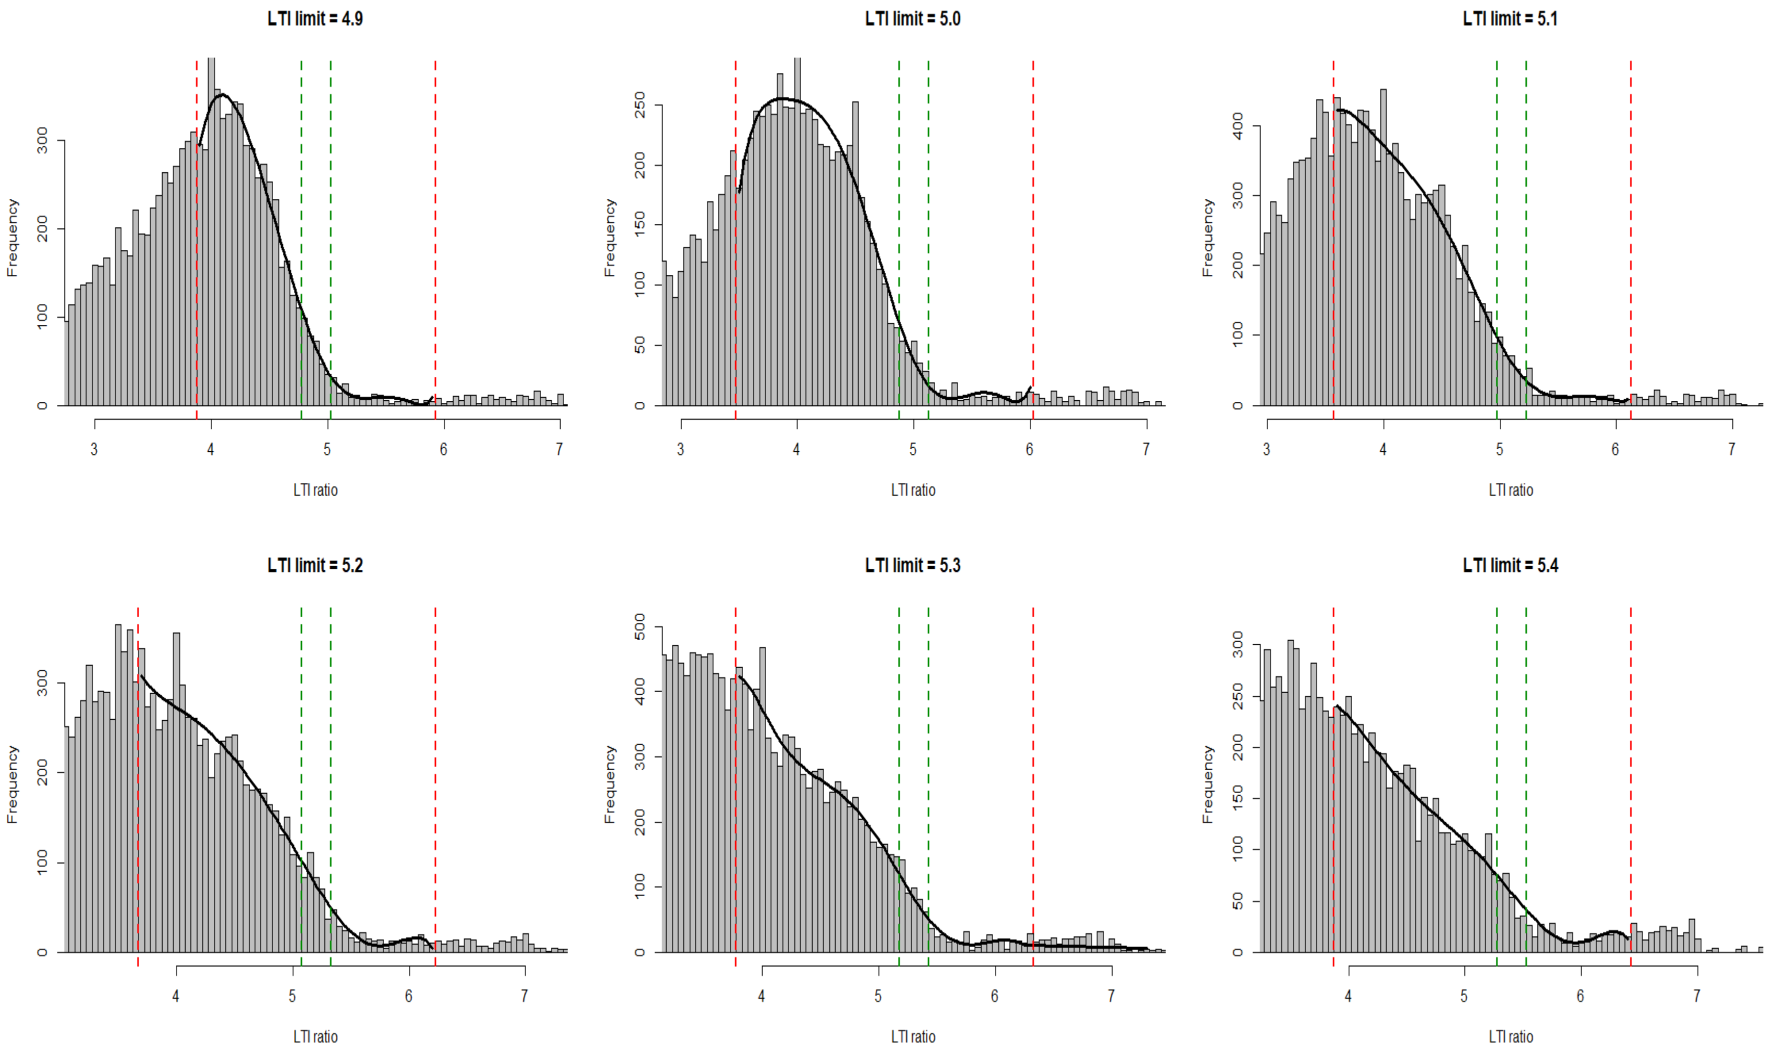

Note: The figure shows the actual LTI distribution (grey bars) of borrowers subject to different LTI limits. Also the figure shows the estimated counterfactual distribution (black line), the LTI limit area (the area between the two dashed green lines) and the analysis area (the area between the two dashed red lines). The estimated bunching mass ($\widehat{B}$) corresponds to the difference between the densities between the actual and the counterfactual distribution, within the LTI area.

2.5.3 Bunching at the LTV limit

The parametric specification in eq. 2.14 does not provide the best approach to evaluate the effect of the LTV limit. The fact that the LTV limit is common to everybody and decreases deterministically by 1% every year implies that the limit cannot explain cross-sectional differences in the amount of debt at origination[25]. Therefore, a bunching approach that evaluates the distributional effects of the LTV rule is better suited. However, also the bunching approach applied in the previous section is not ideal either to investigate the effect of reductions in the LTV limit. The reason is that the LTV rule is much stricter than the LTI rule, meaning that there is not enough density at the right of the observed distribution to obtain a smooth estimate of the counterfactual distribution[26]. In other words, the smoothness assumption on which the previous bunching approach is based would be violated.

To evaluate the effect of the further reductions in the LTV limit, I overcome this issue by using an observed LTV distribution instead of an estimated counterfactual distribution. In particular, I use the restrospective information available in the Loan Level Data to retrieve the LTV distribution of mortgages originated in 2012. The LTV limit reductions were in fact announced in November 2012 and implemented since the beginning of 2013. In this way, I observe how the LTV distribution looked like before the LTV rule was in place, and I use this as a control distribution to assess the effects of further LTV limit reductions.

My approach is similar to the one developed by De Fusco et al (2019) who compare the change in the distributions of originated mortgages in two segments of the U.S. credit market, one of which is exempt to the regulation they study (a rule part of the Dodd-Franck Act). They identify the effect of the regulation using cross-sectional variation in the distributions. Instead, I exploit time variation in the LTV distribution to estimate how changes in the location of the kink in borrowers' budget sets (due to decreases of the

[25] Also, the LTV limit is proportional to house valuations. According to the rule, borrowers are allowed to borrow up to a fraction $(1-\delta)$ of the their property valuation $p_t h_t$. As shown in eq. 2.13, an increase in house prices mechanically translates to an increase in borrowing capacity allowed by the LTV rule. For instance, in our sample period, the LTV limit (δ) has been reduced by 1% every year, but the national house price index has increased by an average rate of 4% a year.

[26] Also, the levels of LTV is among the highest in the world, so it would still be difficult to imagine that banks are willing to grant many loans above the original 106% LTV limit.

LTV limit) affect the LTV distribution and the share of constrained borrowers. In fact, according to Kleven (2016), in case of cross-sectional or time variation in the size or in the location of the kink it is possible to identify the behavioral response as the difference in bunching.

I estimate the effect of LTV limit tightenings in two steps. First, to make the distributions directly comparable, I normalize the total number of loans in each bin by dividing each LTV bin-count by the total number of loans in the analysis area. That is:

$$\overline{n}_j^t = \frac{n_j^t}{\sum_{j=1}^{J} n_j^t} \qquad t = 0, .., T \tag{2.18}$$

Where $\overline{n}_j^t$ denotes the density in bin j of the LTV distribution at time t. Again, $j = 0, .., J$ denotes the bins included in the analysis area[27]. Also, this normalization rules out any extensive margin response, and makes sure that differences in the distributions only reflect responses on the intensive margin, i.e. on the actual Loan-to-Values. Second, I deal with the fact that changes in the distributions over time may be due to reasons unrelated with the change in the LTV regulation. For example, the model in section 2.3 shows that the actual LTV, obtained by dividing the unconstrained mortgage demand in eq. 2.5 by the property value, is an increasing function of house prices. Intuitively, this means that borrowers who bought at the top of the market in 2017 may have been forced to take higher LTV loans than borrowers who bought a house in 2013, when prices were much lower. As a consequence, the distribution skews to the right and the LTV limit is more binding, even without changing the LTV limit. Once again, house price increases can potentially confound the causal effect of the regulation.

To deal with this issue I estimate the following two statistics:

$$\widehat{B_t} = \sum_{j=-k}^{k} \left(\overline{n}_{c+j}^t - \overline{n}_{c+j}^0 \right) \qquad \qquad \widehat{M_t} = \sum_{j=k+1}^{J-c} \left(\overline{n}_{c+j}^t - \overline{n}_{c+j}^0 \right) \tag{2.19}$$

Where $\overline{n}_{c+j}^0$ denotes the density in bin $c + j$ in the LTV distribution at time $t = 0$, i.e. the latest LTV distribution observed before the introduction of the new LTV rule. The density at the LTV limit is the density in a $\pm k$ window[28] around $\overline{n}_c^t$.

[27]We consider bins of size 0.5 in the range [85-110] of the LTV distribution

[28]I set $k = 0.5$ LTV.

The first statistics denotes the Bunching estimate. This is analogous to the one computed in the previous section, but is now is obtained using the 2012 LTV distribution as a counterfactual distribution. The second statistic denotes the Missing Mass to the right of the LTV limit in the LTV distribution at time t , that I use to account for confounding factors such as increasing house prices. The intuition is the following: as the LTV limit decreases, more borrowers get constrained, and the bunching mass increases. Equivalently, as house prices increase many borrowers end up borrowing at higher LTVs, causing the bunching mass to increase too. Over time however, confounding factors can shape the LTV distribution only at the left of the LTV limit, given that under the new rule the limit is binding. On the contrary, changes in the LTV distribution to the right of the (new) LTV limit can only be explained by the progressive LTV limit tightenings. I estimate the Missing Mass to the right exactly to disentangle changes in the bunching mass that are directly attributable to the LTV regulation from those attributable to other factors, such as changes in house prices. This statistic identifies the excess number of borrowers that (i) borrow at the LTV limit and (ii) would have borrowed at higher LTV loans in absence of the regulation.

Results are reported in Table 2.5. In 2013, a few months after the new regulation was announced, the LTV distribution looked very similar to the one computed from mortgages originated one year before. The Bunching, i.e. the estimated excess mass at the new 105% LTV limit was just 2.4% higher than the year before, when the LTV limit was 106%. At the right of the limit, the distributions were also very similar, and the difference between the densities at LTVs higher than 105% (i.e. the Missing Mass) was close to zero, and equal to around -1.5%. As time went by, the LTV limit was progressively reduced, and more and more borrowers got constrained. The Bunching estimate increases from 2.4% in 2013, to 27,5% in 2015 and up to 35.9% in 2018. Correspondingly the Missing Mass, i.e. the part attributable to the policy change, progressively jumps from +1.5% in 2013 to -45.4% in 2017.

This result indicates that progressive reductions of the LTV limit by 5% (from 106% to 101%) contributed to make this constraint binding for up to about 45% of borrowers taking out loans with LTVs in the range between 85 and 110.

Table 2.5: Bunching at the LTV limit

Limit = 105%	est.	95% conf. int.	**Limit = 104%**	est.	95% conf. int.
$\widehat{B}$	0.024***	[0.021 ; 0.030]	$\widehat{B}$	0.212***	[0.208 ; 0.217]
$\widehat{M}$	+0.015***	[0.018;0.012]	$\widehat{M}$	-0.148***	[-0.143 ; -0.152]
N	25.224		N	52.171	
Year	2013		Year	2014	
Limit = 103%	est.	95% conf. int.	**Limit = 102%**	est.	95% conf. int.
$\widehat{B}$	0.275***	[0.270; 0.279]	$\widehat{B}$	0.307***	[0.303 ; 0.311]
$\widehat{M}$	-0.231***	[-0.223 ; -0.236]	$\widehat{M}$	-0.321***	[-0.314 ; -0.326]
N	56.755		N	62.537	
Year	2015		Year	2016	
Limit = 101%	est.	95% conf. int.	**Limit = 100%**	est.	95% conf. int.
$\widehat{B}$	0.330***	[0.328 ; 0.334]	$\widehat{B}$	0.359***	[0.356 ; 0.363]
$\widehat{M}$	-0.401***	[-0.395 ; -0.407]	$\widehat{M}$	-0.454***	[-0.447 ; -0.460]
N	67.795		N	56.557	
Year	2017		Year	2018	

Note: The Table reports LTV bunching estimates in all LTV limits. The Table shows the bunching mass ($\widehat{B}$), the missing mass to the right of the LTV limit ($\widehat{M}$). Also, the Table reports the corresponding 95% bootstrapped confidence intervals. On the bottom of each panel the Table reports the year and total number of starters (N). The symbol *, ** and *** denote conventional statistical significance levels.

Figure 2.8 provides a graphical representation of the result. It shows that the LTV distribution hasn't changed much in the few months between the announcement and the introduction of the LTV rule in 2013. Instead, as of 2014, a considerable amount of borrowers started bunching at the limits. Interestingly, the LTV distribution hasn't changed for values lower than the LTV limit. As the LTV is further tightened, all of the difference observed at the LTV limit is explained by differences at the right part of the distribution, while the distributions look identical to the left. This provides clear support of the second empirical implication in section 2.3.1, represented in Figure 2.3.

Figure 2.6: Bunching at the LTV distribution

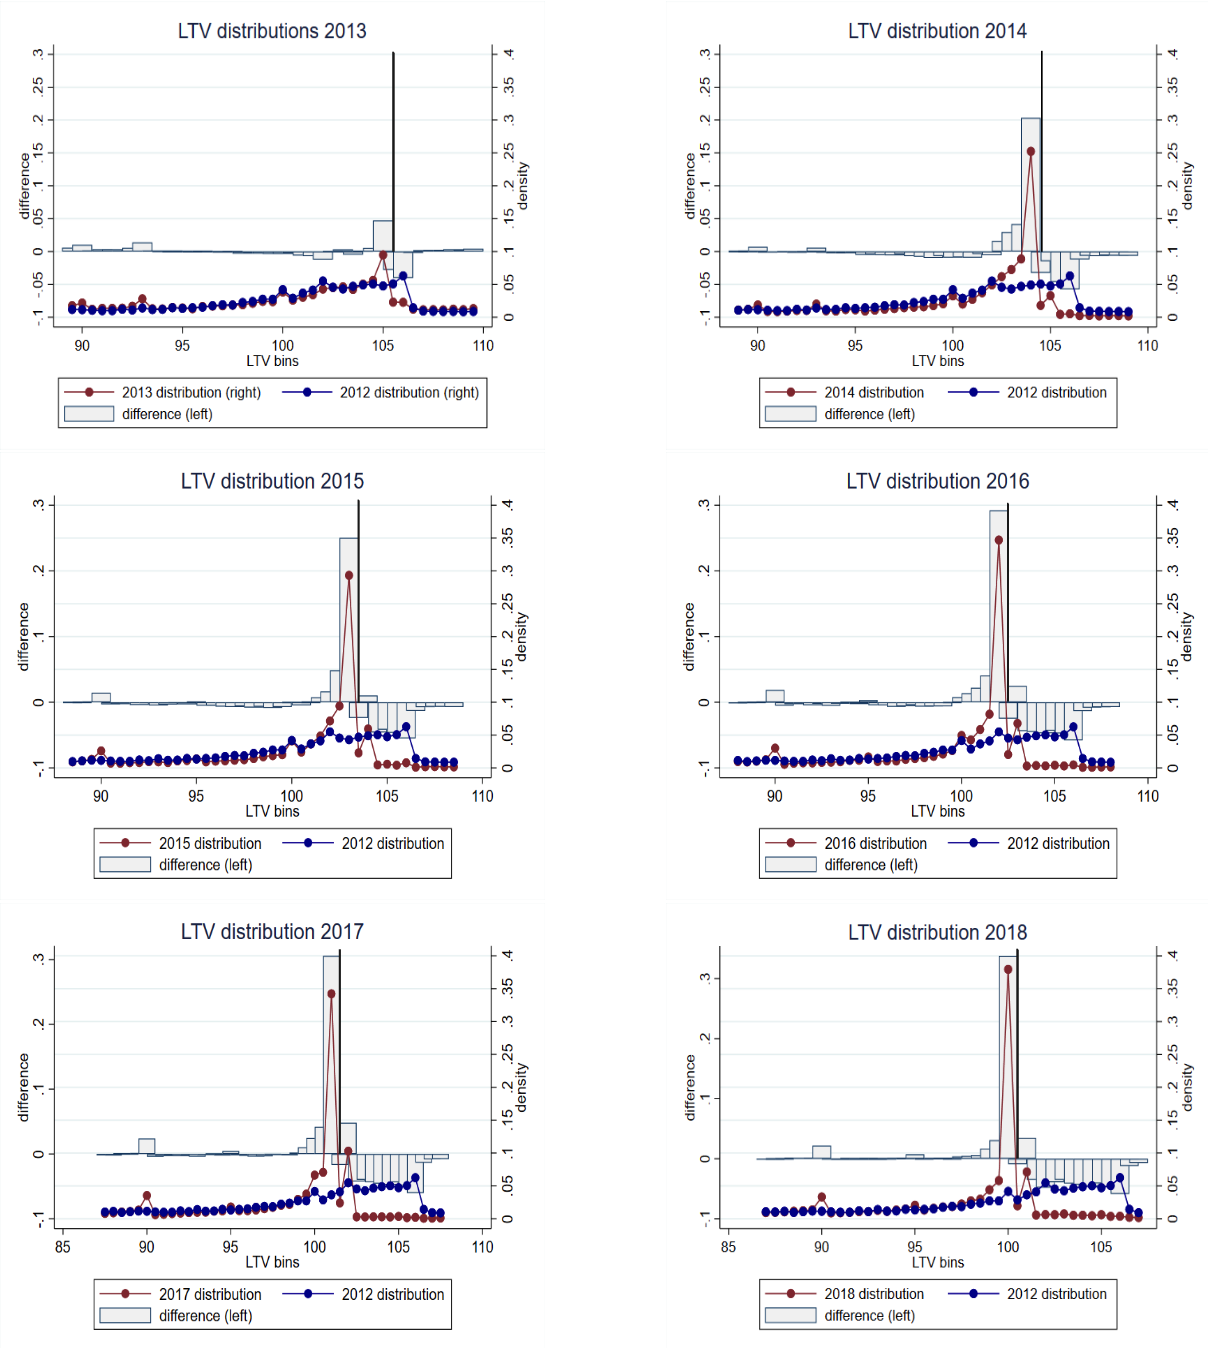

Note: The figure shows the comparison between LTV distributions in each year (red line) and the LTV distribution in 2012 (blue line) in a window of observation bins around the LTV limit, and their within-bin density difference (grey bars, left axis). The black vertical line denotes the LTV limit in place in the current year. The difference between the red and blue densities at the LTV limit (the gray spike next to the vertical black line) represents the estimated Bunching ($\widehat{B}$). The difference between the red and blue densities in each bin at the right of the LTV limit (the gray area at the right of the vertical black line) is the estimated Missing Mass ($\widehat{M}$).

An open question is whether it is possible to attach a causal interpretation to the effect estimated with the Missing Mass $\widehat{M}_t$. Thanks to the use of an observed distribution, instead of an estimated distribution as a counterfactual, this approach does not require to assume local smoothness of the counterfactual distribution around the limit. However, the use a counterfactual distribution observed in the past, implies that the missing mass identifies the effect of the policy only under the assumption that, without the policy, the LTV distribution wouldn't have changed over time. Since the LTV rule applies to everybody, no control group exists and this assumption is not directly testable, as I observe either a distribution subject to an LTV limit, or one not subject to it. To provide evidence in this direction, in Figure 2.7 I compare the LTV distributions of mortgages originated in the four-year before the introduction of the policy, again using the retrospective data available in the LLD, and considering the same LTV range and same scales used before.

Figure 2.7: Pre-treatment LTV distributions

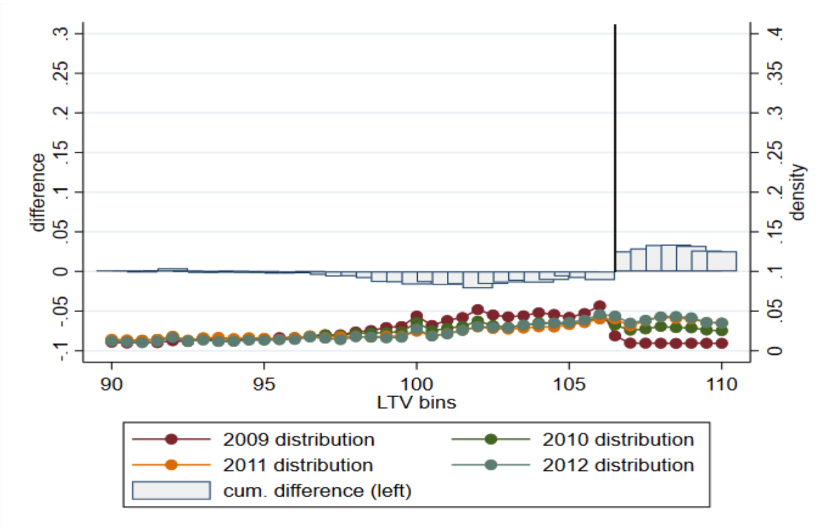

Note: The Figure shows the LTV distributions in the four years before the LTV reductions. The gray bars denote the cumulative difference, i.e. within-bin difference between the densities in the 2012 and 2009 distribution

The figure shows that the LTV distributions didn't change much before the LTV rule and, contrarily to the 2013-2018 period, the colored lines are almost overlapping in the whole support. This evidence supports the validity of the aforementioned assumption. Interestingly, the figure shows that despite the differences were small, the mass at the right of the 106% LTV limit announced (but not yet effective) in 2012 was positive ($\widehat{M}_t = 0.04$).

This means that in the 2009-2012 period banks increased the supply of high LTV mortgages and that, without the introduction of the LTV, they might have engaged in more and more risky lending. This suggests that, if anything, our estimated effect of LTV limit tightenings would more likely be under-estimated rather than over-estimated.

2.6 Concluding remarks

2.6.1 Conclusions

In this paper, I investigates the micro effects macroprudential policies. In particular, I study the effect of leverage constraints in the form of Loan-to-Income (LTI) and Loan-to-Value limits, on household borrowing decisions in a housing market boom.

Using a stylized model of borrowing, I show that increasing house prices induce increasing demand of debt by households, as well as increasing supply of banks via the LTV limit, as increasing collateral values translate into increasing borrowing capacity. The resulting debt level is higher and closer to the binding LTI limit. In the empirical analysis, I investigate this channel and show how effective LTI and LTV limits are in containing household debt, during a period of persistently increasing house prices.

Regarding LTI limits, I find that these have highly heterogeneous effects along the income distribution. For low-income borrowers, changes in the LTI limit induce a changes in household debt of the same size. For high-income borrowers, changes in household debt are decoupled from the limit. Also, I find that the exceptions established by the comply or explain regulation play an important role in giving flexibility to both borrowers and lenders, and likely rule out important effects on the extensive margin, i.e. the possibility that borrowers may not be able to participate to the mortgage market.

Regarding the LTV regulation, I find that the gradual 6 percent reduction in the LTV limit occurred in the 2013-2018 period impacted on up to 45 percent of highly leveraged borrowers. The effects of the LTV regulation are evident in the post-reform LTV distributions, which feature bunching at the leverage limit. The bunching mass at the limit increases as the LTV limit reduces. This suggests that, as the limit is further tightened, more borrowers face reductions in the maximum loan amount they receive, which they must compensate with higher downpayments. The sizable bunching mass at the limit, together with the missing mass at the right of the limit, also show how stricter the LTV regulation is with respect to the LTI regulation.

Regarding house prices, I find that a one standard deviation increase in house prices translates into an increase in household debt by 0.6%. The size of this effect is sizable, and highlights how the housing boom represents an additional constraining factor in households borrowing decisions. I show increasing house prices induce a shift in the level of debt towards the macro-prudential limits for reasons unrelated to changes in the macro-prudential regulation, and if not properly dealt with, this can confound the estimate of the causal effect of the regulation.

2.6.2 Policy implications

I draw two main policy implications from this study. The first policy implication is that to properly control the level of household debt in periods characterized by overheating housing markets, a macro-prudential policy must necessarily impose rules based on households' debt affordability and that are independent on housing market conditions. This is apparently not the case in the vast majority of countries where, according to the dataset on the use of macro-prudential policies made available by Cerruti et al (2017), the macro-prudential policies are mostly based on LTV rules rather than LTI or DSTI rules . Our results show that allowing for household heterogeneity and for flexibility options can significantly reduce the impact that a house price shock can cause.

The second policy implication is that to properly control risk intake in times of increasing risks, a macro-prudential policy should not be pro-cyclical. In good times characterized by decreasing interest rates and increasing house prices, macro-prudential policies often allow borrowers to automatically borrow more. For example, decreasing interest rates induce an automatic increase in the LTI limit, given that these are based on the income and interest rate paid by the borrower. Instead, increasing house prices allow household to increase the amount they can borrow, as the borrowing capacity implied by the LTV rule is proportional to the level of house prices. Therefore, reducing the pro-cyclicality of these rules help in reducing risk-intake in periods of increasing risks.

Chapter 3

Staggered Wages, Unanticipated Shocks and Firms' Adjustments

3.1 Introduction

The extent to which macroeconomic shocks affect employment levels depend on labour market rigidities, which limit the extent to which firms can respond to changing economic conditions with changes in their workforce and/or changes in wages. Labour market rigidities ultimately result from the institutional environment: employment protection legislation prevents firms to cut employment as productivity falls, while minimum wage schemes and downward nominal wage rigidity prevent firms to cut wages during downturns.

This paper evaluates the effect of a particular source of wage rigidity, namely the staggered nature of collective labour agreements, i.e. the phenomenon that wages adjust in a staggered way to changes in the economic environment .

This chapter is taken from a joint project with Jante Parlevliet and Mauro Mastrogiacomo. The views expressed in this article are the ones of the authors and do not necessarily reflect the position of De Nederlandsche Bank. We thank Laurens Harteveld of the employer association AWVN, Jacqueline Twerda, Daphne de Wild and Gijs Lokhorst of the trade union CNV and Martin Schaeps of the Ministry of Social Affairs (SZW) for the useful information provided. Moreover, we would like to thank Egbert Jongen, Jurriaan Paans, Maurice Limmen, Jannie Mooren, Bart van Riel, as well as seminar participants at DNB and CPB for the very useful comments.

Wages are staggered when contract decisions are taken at different points in time and these decisions are valid for a certain number of periods (Taylor, 1979)[1]. At any point in time, only some industries reset their wages, and during the bargaining process employers and unions are likely to respond to events happened in the previous period. Also, as the bargained wage is valid for a number of periods, unions form inflation expectations for the validity period in order to insure workers against real wage losses. Expectations on productivity and profitability are likely to be relevant as well, as unions would make sure that wages grow as much as industry profits grow. Sudden changes in expectations can therefore induce wage dispersion across sectors bargaining in nearby months.

Staggered wages, also known as Taylor contracts, have been introduced in macroeconomics to understand the persistence of inflation dynamics and the transmission of monetary policy[2]. In the most recent macroeconomic literature, staggered wages are the most common method of incorporating labour market rigidities in quantitative macroeconomic models (Taylor, 2016), and represent a nominal friction that helps these models in better reproducing the inflation dynamics observed in the data (Christiano, Eichenbaum and Evans (2005), Smets and Wouters (2007)). However, this wage rigidity has not been extensively studied in microeconomic analyses (Taylor 1979). Notable exceptions include LeBihan et al (2012) and Barattieri et al (2014) who use quarterly data for France and the U.S. to study the hazard rate of wage resets, and Fougère et al (2014) who study industry-level wage floors adjustments and their relation with changes in the minimum wage.

In this paper, we take a microeconomic perspective to empirically assess how staggered wages affect firms' wage versus employment labor cost adjustments. We use a large matched employers-employees dataset obtained from the combination of several administrative data sources to show the effect of contract staggering on a battery of firm-level labour market outcomes. In particular, we consider firms' responses on both the extensive margin (total and flexible employment levels and the number of vacancies in each firm) and the intensive margin (the total amount of bonuses, non-wage benefits, and overtime hours paid to the workers), to understand which adjustment margins firms use the most.

[1] This wage rigidity is different from the previously mentioned downward nominal wage rigidity (DNWR), which refers to the absence of nominal wage cuts. Downward nominal wage rigidity has been studied in Altonji and Devereux (2000), Elsby (2009), Schmitt-Grohé and Uribe (2016).

[2] See Friedman (1977), Fischer (1977), Calvo (1983) and Roberts (1995)

Identifying the effect of nominal rigidities in reduced-form microeconomic analyses is challenging for at lest three reasons. First, in centralized bargaining systems wage negotiations often take place at the industry level, meaning that the identifying source of variation is at a more aggregate level than the unit of analysis. Second, in periods of stable inflation, it is unlikely to observe high wage dispersion across sectors bargaining in nearby periods. Last and perhaps most importantly, the period in which the parties bargain over wage resets may be endogenous, as uncertainty may cause both parties to delay the renegotiation process (Danziger and Neuman (2005)).

To address these identification challenges, this paper focuses on the case of the Netherlands where wage resets are established in collective labour agreements (*collectieve arbeidsovereenkomst*, CAO hereafter) that apply at the sectorial level, but also at the company level for larger firms. The case of the Netherlands is also particularly interesting because nominal rigidities due to contract staggering are very pervasive, due to the large coverage of collective agreements and the relatively long duration of contracts. We then use detailed data on professional forecasts specific to the Dutch economy to identify the timing of a big unanticipated macroeconomic shock that created sufficient wage growth dispersion among CAOs negotiated in nearby months. We identify this shock in October 2008, after which wage growth fell from 3.5 percent to 1 percent in two quarters. Eventually, we exploit a particular feature of the Dutch context to generate plausible exogenous variation around the timing of the shock: in each sector or company in our data, the start date of the subsequent CAO always coincides with the expiry date of the previous CAO. By exploiting variation in the staggered and pre-determined start dates around the aggregate shock, our identification strategy ensures that the starting date of each CAO is independent on the uncertainty generated by the shock. This allows us to estimate the causal effect of nominal wage rigidities on the firms and the workers who signed their CAO before the shock and did not anticipate it.

The paper closest to our own study is Diez-Catalàn and Villanueva (2014)[3] who study the effect of widespread nominal wage rigidity in Spain, where contract staggering is caused

[3]Other related studies are Bertola et al (2012) and all other results from the Wage Dynamic Network, promoted by the European Central Bank, where a total of 17.000 European firms have been surveyed about their wage and price settings, and about the margins of adjustment that they use the most. Other related papers are those of the rent-sharing literature, see Card (2018) and Guiso et al. (2005).

by the presence of province-sector level agreements. They examine workers' total wage growth and transition to unemployment after the Lehman Brothers default, which is an external shock to the Spanish economy, and find that this nominal rigidity increases the probability of transition to unemployment, especially among minimum wage workers. Differently from them, in this paper we take the firm perspective to investigate firms' labour cost adjustments after the start of the 2009 recession in the Netherlands. We improve on this study along two main dimensions: first, using data on collective agreements in place in the Netherlands in the period 2005-2012, we provide evidence on how bargaining has changed during the crisis period. Then, using an identification strategy based on CAOs pre-determined start dates (instead of the endogenous agreement dates), we identify the *causal* effect of nominal rigidities on firms' labour market outcomes.

Our main result is a precisely estimated zero effect of contract staggering on employment. Firms that agreed on a high wage growth before the recession did not cut employment levels more than those firms that were able to anticipate the shock and to adjust contractual wage growth. Instead, we show that non-anticipating firms adjusted labour cost mostly on the intensive margin by cutting the total amount of bonuses and benefits paid to the workers. We show that these cuts fully compensate the higher contractual wage growth. This result partially contrasts with Hall (2005) and Gertler and Trigari (2009) who show that including staggered multi-period bargaining in search and matching models of the labour market, allows to obtain more volatile responses to an aggregate shock of all labour market indicators, including employment. Consistently with these models, we find that the effect is increasing in the degree of rigidity. We show a significant negative effect on employment, other than a stronger negative effect on bonuses and benefits, in sectors covered by collective agreements with contract durations much higher than what is normally assumed in macroeconomic models for the U.S. economy such as Christiano, Eichenbaum and Evans (2005), Smets and Wouters (2007), Gertler and Trigari (2009), where wage resets often occur every three to four quarters. Our results undermine the role of contract staggering in amplifying labour market fluctuations and aggravating employment losses after a shock. At the same time, our results provide clear evidence that labour market rigidities (especially excessively long contract durations) induce firms to use the available adjustment margins to cut labour costs in recession times. Our paper has important implications for the role of collective bargaining. Out of concern for job losses, international organizations have often advocated to reduce nominal rigidities in wage setting, for instance by moving bargaining to the firm level. Yet, such reforms have proven contentious and have raised concerns about the overreliance on wage moderation, the erosion of collective bargaining and the consequences on wage inequality (see Dustmann et al., 2009; Blanchard et al. 2014; Boeri and Jimeno, 2015; Addison et al 2017). Our paper shows that the nominal rigitidies that often result from collective bargaining do not necessarily come at the cost of employment losses, and suggests that building more room for discretionary pay components in collective agreements may be a way to alleviate such rigidities.

The paper proceeds as follows. Section 2 discusses the institutional framework. Section 3 describes the data and the descriptive statistics. Section 4 presents the identification strategy being used and the results of the empirical analysis. In Section 5 we present robustness checks and heterogeneity analysis. Section 6 concludes.

3.2 Institutional Framework

Wage setting in the Netherlands is characterized by a dominant role for collective bargaining. According to the OECD[4] almost 80 percent of all wage earners in the Netherlands are covered by a collective agreement (CAO), as compared to an average of 32 percent across OECD countries. The high coverage is due to the presence of the so-called *erga omnes* provisions, i.e. agreements that automatically apply to all workers within firms that are members in a signing employer association, as well as government-issued extensions of collective agreements to all firms and workers in a sector (Visser, 2016).

The great majority of workers are covered by a CAO that applies to the whole sector, although especially large firms can have their own firm-level agreements. As in many other countries, in the Netherlands for some decades now there has been a shifting towards a more decentralized wage setting. The main route to decentralization has been the increasing scope for customization within sectoral agreements. Examples include the

[4]See https://stats.oecd.org/Index.aspx?DataSetCode=CBC

use of minimum standards that can be topped-up at the firm level and the so-called *à la carte* provisions which allow the bargaining parties of each firm to select and include the wage and non-wage components they prefer. Wage bargaining in the Netherlands is furthermore set in a mature corporatist setting where various coordination mechanisms are in place to sustain collective bargaining: to ensure that confederations have a mandate, unions and employers confederations participate to an annual discussion aimed at setting the guidelines for the next wage increases and other bargaining priorities. This internal coordination is quite strong, especially for sectoral unions that often agree on a maximum wage demand and can possibly team up with employers against dissident unions.

For the purpose of this paper, it is important to mention details on the timing, the frequency and the synchronization of Dutch collective agreements. In theory, wages are bargained according to a specific calendar: based on the forecasts published by the Bureau of Economic Policy Analysis (CPB) in September, parties define their wage demands for the following year and set bilateral guidelines for wage increases and other bargaining priorities. In practice, as we will document in the next section, bargaining takes place throughout the year. This is different from the case of Japan or the U.S. where wage bargaining is synchronized across sectors (Olivei and Tenreyro, 2010). Furthermore, while collective agreements in countries such as Japan or the U.S. typically have a duration of 12 months, CAOs in the Netherlands can have longer as well as shorter durations. On average, the duration of a collective agreement is about 20 months (Visser, 2019), while the average duration of wages among the 17.000 European firms surveyed for the Wage Dynamic Network initiative was about 15 months (ECB, 2009). Hence, the Dutch labour market is characterized by a relatively high degree of contract staggering.

Concerning the validity period of CAOs, it is important to mention that there can be a substantial difference between the start date and the signature date of collective agreements. Like in other countries, the Dutch practice is that in case of no new agreement the old agreements remains valid after expiration (so called *ultra-activity*) whereas in case a contract is agreed after the pre-determined start date, the resulting wage increase can apply retrospectively (*retro-activity*) (Hijzen et al. 2019).

Last, it is worth mentioning that collective agreements typically rule most components of workers' wage bill, such as the ordinary wage, other pay components (e.g. the thirteenth month), contribution to pension funds and benefits (e.g. travel costs). However, firms may have also some room for discretion, such as on individual and collective performance pays, which depend on the performance of the firm.

3.3 Data and Descriptive Statistics

3.3.1 Collective Agreements Data

Our dataset on collective labour agreements (CAOs) is obtained from the Ministry of Social Affairs (MvSZ). The dataset covers most of the CAOs in place in the 2006-2012 period, and contains information on the agreement, the start and the expiration dates of agreements, the contractual wage increase as well as incidental and structural adjustments in other pay components. Also, it contains the Standard Business Classification (SBI) code that identifies the sector of the economic activity of each firms and is used to merge the CAO information with the companies and workers information available in the Matched Employers Employees Data.

Table 3.1: Descriptive Statistics (CAO Data)

	By CAO Start Date				
	2007	2008	2009	2010	2011
Contract duration	19.4	21.4	18.6	18.6	21.7
Agreement delay	2.3	2.0	4.8	2.4	2.9
Wage Growth %	2.9	3.0	1.3	1.2	1.7
N° agreements	88	71	74	73	60
	By CAO Agreement Date				
	2007	2008	2009	2010	2011
Contract duration	19.7	21.0	18.0	18.7	21.2
Agreement delay	2.3	2.2	2.8	3.8	3.2
Wage Growth %	2.9	3.1	1.3	1.1	1.8
N° agreements	95	70	56	80	60

Note: The table reports means and the number of collective agreements by CAO start date (panel above) and by CAO agreement date (panel below). The percentage wage growth is annualized (to normalize over contract duration) and also includes other structural pay increases (e.g. 13th month).

Table 3.1 reports descriptive statistics regarding contract durations, agreement delays and contractual wage growth, and shows how the main bargaining characteristics evolved over time. The duration of collective agreements is the characteristic that has remained most stable. This is due to the fact that in many sectors social partners tend to stick with the same contract duration. Instead, contractual wage growth is more volatile and is closely related to the business cycle: the annualized wage increase and the total annual wage growth (that includes incidental and structural adjustments in other pay components) were around 3 percent in the pre-crisis period, then fell to 1.3 percent during the 2009 crisis and slightly recovered to 1.8 percent in 2011. More importantly, the table shows that on average collective agreements are signed with some delay. Figure 3.1 shows the distribution of the agreement delay before, during and after the 2009 recession.

Figure 3.1: Agreement Delay

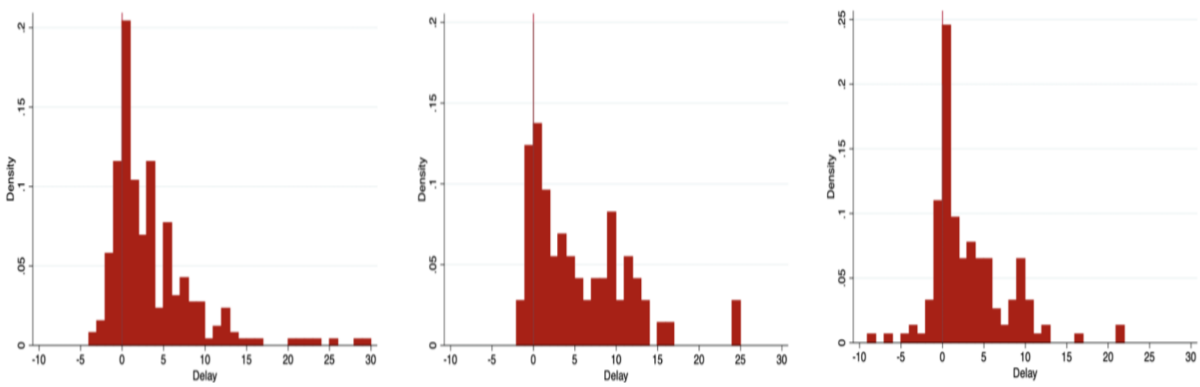

Note: The figure shows the distribution of the agreement delay of contracts started before (left), in (centre) and after (right) 2009. The agreement delay is expressed in months.

The figure shows that social partners typically find an agreement in the months around the expiry date of the previous contract (which corresponds the start date of the new agreement). Around 25% of contracts are agreed exactly on time, while in many other cases agreements are reached either a few months before the start date, or with a delay of one to several months. For the purpose of this paper, it is interesting to notice that agreement delays have increased, on average, during the crisis period: in 2009 (middle panel) the share of contracts agreed on time has dropped by more than 5% and the overall distribution is more skewed to the right, as compared to the distributions before and after 2009. Instead, the distribution of the agreement delay before and after 2009 are very similar to each other. This descriptive evidence supports the idea that agreement dates are endogenous. The parties involved in the bargaining process may have decided to delay their decisions or postpone the whole negotiation process when the shock hit the economy. To study the consequences of contract staggering in the recession period, the empirical analysis focuses on the analysis of reference CAOs. We define the reference CAO as the collective labour agreement that was in place during the 2009 recession in each company and sector[5]. Table 3.2 reports the list of all reference CAOs and shows the agreed percentage wage increase and the validity period of each reference CAO. The table shows that among all collective labour agreements in place during the recession, about 25% of them consist of company-specific CAOs. Also, it shows that the agreed percentage wage increases in these reference CAOs range from 0 to 3.5 percent, and that the length of the validity period of a CAO ranges from six months to about five years.

Table 3.2: CAO data

CAO name	Wage increase %	Validity period	CAO name	Wage increase %	Validity Period	CAO name	Wage increase %	Validity period
ABN-AMRO	3.50	Jan-09 Mar-10	Gemeentelijk vervoersbedrijf (GVB)	-	Jan-09 Jan-10	Postkantoren	1.00	Apr-09 Jan-10
AGF	0.00	Apr-09 Apr-10	Gemengde- en speelgoedbranche	1.25	Jan-09 Jan-11	Praxis group	0.29	Feb-09 Feb-11
ANWB	3.25	Apr-08 Apr-10	Getronics Nederland	0.00	Jan-09 Jan-10	Primair Onderwijs	0.84	Jan-09 Jan-10
Achmea	3.25	Jan-08 Jun-09	Glastuinbouw	2.5	Jul-07 Jul-10	Provinciepersoneel	0.70	Jun-07 Jun-09
Afbouw	2.00	Jan-08 Jan-10	Goederenvervoer Nederland (KNV)	3.50	Oct-08 Jan-10	RABO-bank	3.00	May-07 May-09
Akzo Nobel Pharma bv	3.00	Apr-07 Apr-09	Grafi-media	2.75	Feb-07 Feb-10	Recreatie	3.50	Jul-08 Jul-09
Albert Heijn (distributie org.)	3.50	Oct-08 Oct-09	Heineken Nederland Beheer BV	2.50	Jul-08 Oct-09	Reisbranche	1.00	Apr-09 Apr-11
Algemene Banken	3.00	Jan-08 Apr-10	Hema	0.75	Feb-09 Feb-10	Rijkspersoneel	2.30	Jan-07 Jan-11
Apotheken	3.25	Apr-08 Apr-10	Hoger beroepsonderwijs (HBO)	1.30	Aug-07 Aug-10	SNS Reaal	3.25	Jun-08 Jun-09
Arbo Unie Nederland	1.00	Jul-08 Jul-11	Holland Casino	0.00	Apr-09 Jan-11	SNT Nederland	1.13	Jul-08 May-10
Architectenbureaus	1.00	Apr-09 Apr-10	Horeca- en aanverwant bedrijf	2.76	Apr-08 Apr-10	Schilders- en glaszetbedrijf	1.50	Mar-09 Mar-11
Atos Origin	4.00	Jan-08 Jan-09	Hoveniersbedrijf	0.60	Mar-09 Mar-10	Schoonmaak	3.50	Apr-08 Jan-10
Bakkersbedrijf	3.50	Jun-08 Aug-09	Huisartsenzorg	1.50	Jan-09 Apr-10	Selektvracht B.V.	1.50	Apr-09 Jan-11
Beroepsgoederenvervoer	3.50	Oct-08 Jan-10	IKEA Nederland b.v.	3.25	Oct-08 Oct-09	Slagersbedrijf	1.00	Apr-09 Apr-11
Beroepsonderwijs	0.07	Feb-09 Apr-11	ING	3.00	Jul-08 Jan-12	Sociale Verzekeringsbank	3.25	Apr-08 Apr-10
Betonproduktenindustrie	1.00	Mar-09 Apr-11	Informatie & communicatie	3.00	Jan-08 Jan-10	Sociale werkvoorziening	3.00	Nov-07 Mar-10
Beveiligingsorganisaties	3.25	Aug-08 Jul-11	Jeugdzorg	3.50	May-08 May-10	TNT N.V.	0.70	Apr-09 Jan-12
Bloemen en planten (groothandel)	3.00	Jan-08 Jul-09	KLM-grondpersoneel	1.25	Apr-09 Apr-10	TNT Post B.V.	2.86	Apr-08 Apr-09
Bloemen- en Plantendetailhandel	1.40	Jan-09 Apr-11	KPN Contact	2.00	Jul-08 Jan-10	Tankstations en wasbedrijven	0.00	Jan-09 Jan-10
Boekhandel en Kantoorvakhandel	1.75	Jan-09 Feb-10	KPN N.V.	0.00	Jan-08 Jan-10	Taxivervoer	1.50	Mar-09 Jan-14
Bouwnijverheid	1.75	Apr-07 Jul-09	Kappersbedrijf	1.00	Jan-09 Jul-10	Technische groothandel	3.25	Jan-08 Apr-10
Canon	3.50	Apr-08 Jul-09	Kinderopvang	2.00	May-08 May-09	Technische installatiebedrijven	3.50	Feb-08 Dec-09
Carrosseriebedrijf (metaal)	3.50	Feb-08 Dec-09	Landbouwwerktuigen exploiterende	2.00	Apr-07 Apr-09	Timmerindustrie	1.25	Jan-09 Jul-09
Contract-cateringbedrijf	3.50	Apr-08 Apr-10	Levensmiddelen (groot.)	3.25	Apr-08 Apr-10	UWV (Uitvoeringsorgaan	2.60	May-07 May-10
Corus Staal B.V.	0.50	Apr-09 Apr-10	Levensmiddelen en/of zoetwaren	2.00	Jul-07 Jul-09	Uitzendbureaus	1.00	Jan-09 Apr-11
DHL Express Nederland	1.50	Apr-09 Jan-11	Levensmiddelenbedrijf	3.25	Apr-08 Apr-10	Unilever	2.00	Mar-09 May-10
DSM Limburg B.V.	3.50	Jun-08 Jun-09	Metaalbewerking (metaal)	3.50	Feb-08 Dec-09	Universitair Medische Centra	2.10	Jan-08 Mar-11
Defensie-personeel	1.00	Mar-09 Mar-10	Metalektro	3.00	Nov-07 Feb-10	Vroom Dreesman	2.05	Feb-08 Feb-10
Delta Lloyd n.v.	3.00	Jun-07 Jun-09	meubileringsbedrijven	1.15	Jul-08 Jul-10	VVT	3.25	Jan-08 Mar-10
Dierhouderij	2.50	Jul-07 Jan-10	Mode- en sportdetailhandel	2.00	Jan-08 Jul-10	Verzekeringsbedrijf (binnendienst)	3.00	Jun-07 Dec-09
Doe het zelf branche	2.34	Jan-07 Jul-09	Mode-; interieur-; tapijt- en	3.50	May-08 Jul-10	Vleessector	1.00	Apr-09 Apr-11
Drogisterijbranche	1.00	Apr-09 Oct-10	Motorvoertuigen	3.50	Feb-08 Jan-10	Voortgezet Onderwijs	3.00	Jul-08 Aug-10
Electrotechnische detailhandel	1.58	Jan-09 Jan-12	Nederlandse Spoorwegen (nieuw)	3.00	Apr-07 May-09	Vroom en Dreesmann Food	1.92	Feb-08 Feb-10
Energie	2.80	Jan-08 Jul-09	Oce Nederland B.V.	3.50	Apr-08 Jul-09	Welzijn en maatschappelijke dienst	2.50	May-08 Jan-12
Facilitaire Contactcenters	2.25	May-08 May-10	Ons Middelbaar Onderwijs	3.00	Jul-08 Aug-10	Wonen	2.75	Jan-08 Jan-10
Fortis-bank	3.50	Jan-09 Mar-10	Open Teelten	1.00	Jan-09 Jul-10	Woondiensten	1.50	Jan-09 Jan-11
Geestelijke Gezondheidszorg (GGZ)	1.00	Apr-09 Mar-11	Openbaar vervoer	3.50	Jan-08 Jul-09	Ziekenhuizen	1.00	Mar-09 Mar-11
Gehandicaptenzorg	1.50	Jan-09 Mar-11	Philips (nieuw)	-	Jan-09 Jan-10	Zoetwarenindustrie	3.00	Jul-07 Jul-09
Gemeente-ambtenaren	2.20	Jun-07 Jun-09	Politie-personeel	3.50	Jan-08 Jan-12	Zorgverzekeraars	3.00	Jun-07 Jun-09

Note: The table shows the list of all reference collective labor agreements (reference CAOs) in our sample and reports the annualized agreed percentage wage increase and the validity period of the agreement.

[5]In case of two CAOs in place during these period, the reference contract has been defined on the basis of the number of months in 2009 in which each of them was in place or on the basis of information availability.

3.3.2 Matched Employers Employees Data

The data that we use is a matched employers employees dataset obtained by combining different administrative sources available from Statistics Netherlands. For the employers part, we combine information available in the General Firm Register on the size, the municipality and the sector in which the firms operate with balance sheet information coming from the corporate income tax data of the tax authority. For the employees part, we combine contract spells and monthly wage bills from the tax authority with workers characteristics available from the Employee Insurance Agency (UWV) and the General Civil Register. The Matched Employers Employees Dataset is then obtained by merging the employers data with the employees data via the firms and workers identifiers. Eventually, the final dataset is obtained by further merging the data on collective labour agreements merged via the SBI codes[6].

Table 3.3: Descriptive statistics (MEED)

	Panel A: Employers				**Panel B: Employees**		
	2007	2008	2009		2007	2008	2009
Sales (in '000s)	5.768,70	5.448,44	4.561,50	Net Annual Wage	21.271,1	22.074,4	23.164,3
	(142.448,1)	(129.001,9)	(104.524,4)		(31.001,1)	(24.827,5)	(24.440,0)
Total Assets (in '000s)	5.497,33	6.292,45	5.890,95	N° hours worked	23,4	23,6	23.8
	(325.994,5)	(322.626,7)	(312.336,7)		(14,1)	(14,2)	(14,0)
Net Profits (in '000s)	652,19	455,79	255,17	Bonus	2.120,2	2.303,1	2.525,1
	(63.938,7)	(55.10267)	(21.648,2)		(15.105,0)	(10.813,9)	(9.963,0)
Wages (in '000s)	805,90	794,21	730,35	Bonus share (%)	13,8	13,7	13,2
	(12.943,6)	(14.320,3)	(12.962,7)				
N° of Employees	38,7	39,8	38,8	Bonus share (interquantile range)	6,3	7,1	7,4
	(572,8)	(634,5)	(582,0)				
New workers (Inflows)	18,8	16,9	14,1				
	(57,1)	(60,0)	(43,8)				

Note: The table shows descriptive statistics of the Matched Employers Employees Dataset. Descriptive statistics of the employers data are reported in the left panel (Panel A), while descriptive statistics of the employees data are reported in the right panel (Panel B). The table reports means and standard deviation (in brackets) of each reported variable.

Table 3.3 reports descriptive statistics of the Matched Employers Employees Data (MEED). It contains the most important attributes of both the employers and employees side of the MEED. In panel A we report summary statistics of some firm-level attributes, while in panel B we report summary statistics for the main worker-level attributes.

[6]In case of company-specific CAOs, we combine this information on the basis of a statistical match based on the SBI code, the municipality where the company's headquarter is based, and the total number of employees.

The employers data contains information for a total of 90.688 firms that have at least five employees and that have been successfully matched with their corresponding employees data[7]. These firms are also observed in the whole sample period which spans from 2007 to 2009. For all these firms we observe sector-level profitability indicators and aggregate financial statistics linked via the sector codes, while for a 20% subset of firms in the manufacturing sector we observe the actual balance sheet data merged via the employer identifier. For listed companies, that typically have a company specific collective agreement, we obtain the corresponding balance sheet information from Bloomberg.

The employees data contains information on the population of workers in the Netherlands. In particular, we observe contract spells and monthly wages of all workers of the selected 90.688 firms. Thanks to the available employer identifiers, contract spells data allows us to observe, at any point in time, the total employment levels of each company included in the sample. The total number of workers employed in our sample of firms is equal to 5.633.602, 1.541.430 of which were employed in the same firm for the whole period. The monthly wage bills data allows us to observe the whole structure of workers' wages, and allows us to obtain a decomposition of their remuneration into the following components: the ordinary wage, the bonus and benefit components, and the part attributable to taxes and contributions[8]. In this paper, we focus mostly on the bonus components. As shown in Table 3.3, this represents a fair share of workers' remuneration. Eventually, we set the data at the firm level and at quarterly frequency to investigate short term firms' labour cost adjustments in the presence of nominal rigidities.

[7]The General Firm Register unfortunately contains two different and non-overlapping identifiers, namely the encrypted fiscal number and the firm identifier, and it is possible to match the employers with the employee data only on the basis of the firm identifier. For this reason, our employers data do not contain the population of firms in the Netherlands. The share of successfully matched companies is about 75% of the total number of firms with at least five employees.

[8]This last component includes health insurance, disability insurance, contributions to pension funds.

3.4 Empirical Analysis

3.4.1 Identification Strategy

The aim of this section is to empirically assess how an aggregate shock propagates via nominal rigidities in the labour market and affects firms' adjustment on the intensive and extensive margin. To address this question, our identification strategy exploits two key elements: the occurrence of an unanticipated shock and exogenous variation in nominal wage rigidities.

The shock that we exploit is the 2009 recession. In terms of GDP, it was the strongest shock to the Dutch economy of the last 40 years and led to a substantial change in expectations[9]. To identify the timing of the shock in such a way that it was unanticipated, we use professional forecasts data on expectations about consumer prices, industrial production, investments and private consumption[10]. Table 3.4 shows that as of October 2008, professional forecasters were still expecting a consumption, inflation and investments growth in the current calendar year, and their one-year ahead forecasts were stable. On the contrary, the recession scenario was already included in their April 2009 forecasts: industrial production and real investments were expected to fall by almost 10%, private consumption was also expected to drop and inflation expectations had been substantially cut. These forecasts are perfectly in line with those published by the Bureau of Economic Policy (CPB) that serve as a common background for the parties involved in the bargaining process[11]. As evident, the change in expectations took place between October 2008 and April 2009. Therefore, to make sure that the shock was unanticipated, we set the timing of the shock in October 2008.

[9] A change in expectations is a necessary condition to produce wage dispersion in a staggered wage setting, given that unions and employers. Gertler and Trigari (2009) show that when contracts are staggered, the bargained wage not only depends on relative bargaining power, but also on expected future economic conditions. Similarly, Hall (2005) shows that changes in the economic environment shift the boundaries of the bargaining set and changes the employers' incentives to recruit.

[10] Consensus data consists of forecasts data at semiannual frequency, specific for individual countries, obtained by surveying and combining opinions of professional forecasters such as advisors, institutional investors and rating agencies. For more details, see https://www.consensuseconomics.com

[11] These can be found in the Macroeconomic Outlook (MEV) published in 2008, available at https://www.cpb.nl/en/publication/macro-economic-outlook-2008, and in the Central Economic Plan (CEP) published in March 2009 and available at https://www.cpb.nl/en/publication/central-economic-plan-cep-2009.

Table 3.4: Descriptive statistics (Consensus data)

	Industrial Production					**Consumer Prices**			
Forecast Horizon	Current year	1 year ahead	2 years ahead	mid-term (5 years)	Forecast Horizon	Current year	1 year ahead	2 years ahead	mid-term (5 years)
October 2007	3.6	2.0	2.0	0.8	October 2007	1.7	2.2	2.3	2.0
April 2008	3.8	1.6	1.7	2.6	April 2008	2.4	2.5	2.0	1.9
October 2008	0.6	0.1	1.0	3.3	October 2008	2.6	2.0	1.7	2.3
April 2009	-9.9	1.0	2.8	5.0	April 2009	1.0	1.0	1.0	2.1
October 2009	-9.1	1.6	2.3	3.1	October 2009	0.9	1.0	1.3	1.6
April 2010	3.7	2.0	0.3	-0.5	April 2010	1.1	2.3	1.7	1.6

	Real Gross Investments					**Private Consumption**			
Forecast Horizon	Current year	1 year ahead	2 years ahead	mid-term (5 years)	Forecast Horizon	Current year	1 year ahead	2 years ahead	mid-term (5 years)
October 2007	4.4	3.4	2.9	1.4	October 2007	2.0	1.9	1.9	1.7
April 2008	3.6	2.0	2.7	2.6	April 2008	2.0	1.7	2.1	1.9
October 2008	5.7	0.0	2.6	3.1	October 2008	1.9	0.7	1.5	2.1
April 2009	-9.3	-2.7	2.0	3.9	April 2009	-0.7	-0.4	1.0	2.1
October 2009	-11.2	-2.8	1.8	2.5	October 2009	-2.7	-0.1	0.7	1.6
April 2010	-3.5	2.4	1.4	2.0	April 2010	-0.1	1.1	1.1	1.4

Note: The table reports, for each date and forecast horizon, the annual percentage expected growth rate of Industrial Production, Consumer Prices, Real Gross Investments and Private Consumption. Each column and row refer to a different forecast horizon and survey date, respectively. Source: Consensus data.

The exogenous variation that we exploit is the variation in the start dates of collective agreements in place during the 2009 recession (reference CAOs). In fact, the start date of a CAO always coincides with the end date of the previous CAO, meaning that once a collective agreement is signed, the start date of the next collective agreement is known. This in turn implies that the validity period of all CAOs was entirely pre-determined when the 2009 recession hit the Dutch economy, and this gives us plausible exogenous variation in the CAO start dates around the timing of the shock.

To test whether the 2009 recession was sudden and strong enough to create wage dispersion, we regress the agreed wage increase on a dummy variable equal to one if the start date of the CAO is before October 2008. The associated coefficient identifies the wage dispersion among CAOs signed around the timing of the shock. Table 3.5 shows the the unconditional difference in the wage increase is 2.79%, which is extremely high given that the total absolute range of the values of this variable is 3.5%, as shown in Table 3.2. Controlling for differences in profitability across sectors only slightly attenuates the difference to 2.50%. Alternatively considering the annualized wage increase (to normalize over the contract duration) or the total annual increase (that further includes allowances) does not change the result.

Table 3.5: Wage Dispersion Test

	Dependent variable:			
	Annualised Wage Increase %		**Total Annual Wage Increase %**	
	(a)	(b)	(a)	(b)
CAO started before 2008Q4	2.78*** (0.156)	2.45*** (0.168)	2.85*** (0.160)	2.35*** (0.163)
Controls	×	✓	×	✓
N	101	101	101	101

Note: The table reports the result of the test for the CAOs' wage dispersion around the 2009 recession. The estimated coefficient identifies the average wage increase differential before and after the beginning of the crisis (difference in means). The first column reports the unconditional difference in wage growth, while the second column shows the same difference after having controlled for the sectorial Return on Assets (ROA) and Return on Equity (ROE), and the sectorial profit margin. The symbols *, ** and *** denote conventional statistical significance levels. Source: Collective agreements data.

This observed wage dispersion across sectors resetting wages around the start of the recession is the key of our identification strategy. The channel is the following: firms covered by a CAO that started before the 2009 recession did not anticipate the shock and committed to pay a high wage increase in a period of decreasing sales. Therefore, they may have incentives to adjust labour cost using the available adjustment margins: they might reduce the number of employed workers by not opening any new vacancy and/or by not rolling over flexible contracts. Alternatively, they might compensate the increase in the contractual wage established in the CAO with a cut in the discretionary part of the bonus and benefit components[12].Let $T_s = \mathbf{1}[\text{CAO started before 2008Q4}]$ be the treatment group indicator. We estimate the following equation:

$$y_{j,s,t} = \lambda_t + c_j + \sum_{k=2008Q4}^{2009Q4} \beta_k (T_s \lambda_k) + \theta\, newCAO_{s,t} + \mathbf{Z}'_{s,t}\gamma + \epsilon_{j,s,t} \tag{3.1}$$

[12]Benefits include travel cost reimbursements, the rental value of any car provided by employer, the rental value of a service-house provided by the employer, and bonuses in the form of real goods. Example of this last category may include holidays or event tickets, which are classified as benefits (instead of bonuses) because they are non-monetary compensations given to the workers.

where $y_{j,s,t}$ is the labour market outcome of firm j in sector s in quarter t, $\mathbf{Z}'_{s,t}$ is a set of control variables and c_j and λ_t are firm and time fixed effects. The effect of interest is captured by the vector of coefficients β_k which identify the Average Treatment Effect on the Treated (ATT): $\beta_k = E(y_{j,s,t}(1) - y_{j,s,t}(0)|s, t = k, T = 1)$. The treatment period lasts from 2008Q4 to 2009Q4. The variable $newCAO_{s,t}$ equals one in all periods in which the reference CAO is replaced by a new collective agreement.

Figure 3.2: Identification Strategy

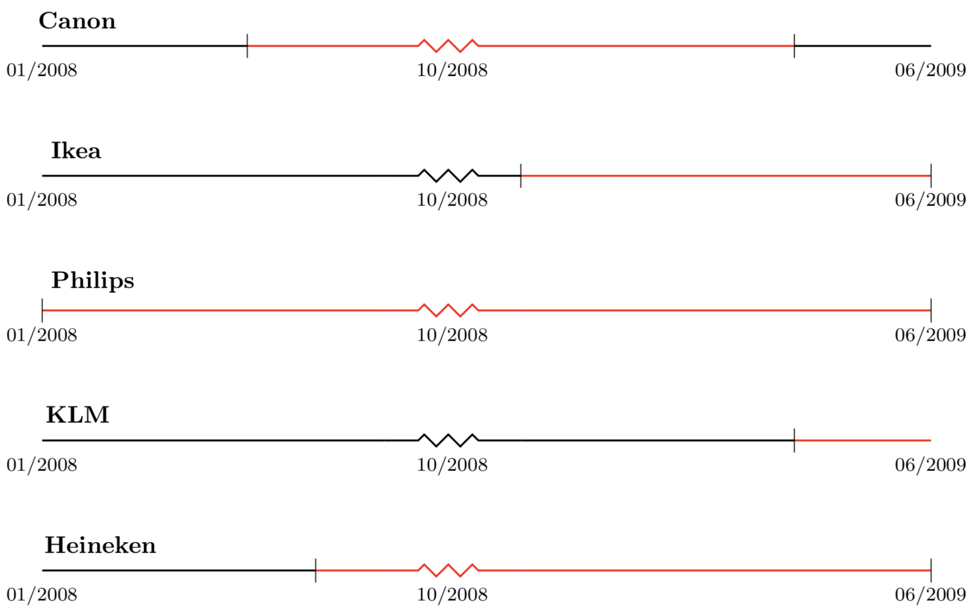

Note: Graphical example of the identification strategy used in the empirical analysis. The figure shows, in a time line, the period in which the reference CAO was in place in five well-known companies. The red segment denotes the validity period of the Reference CAO and the *zigzag* line denotes the timing of the shock. The treatment (control) group consist of firms covered by a CAO that started before (after) October 2008. Here, Canon, Philipps and Heineken are in the treatment group, Ikea and KLM are in the control group. The treatment period is the period between the (common) shock and the (specific) end date of the reference CAO: in this graphic example, Canon early-exits from treatment.

Figure 3.2 shows a graphic example of the identification strategy we use. The variable $newCAO_{s,t}$ accounts for the fact that, es evident from the figure, some CAOs signed in 2008 are already re-negotiated in 2009. Thus, it takes into account for the fact that once a new collective agreement is agreed, the contractual wage growth is adjusted to the new economic scenario. As a result, the treatment period consists of the crisis period, up to the time in which the reference CAO is still in place. Again, the treatment assignment is based on the start date of the reference CAO, which is already established when the previous CAO was agreed, and is thus independent on the uncertainty caused by the crisis.

3.4.2 Results

Table 3.6 reports the result of the analysis of firms' extensive margin responses to the unanticipated shock, and reports the estimated ATTs and their cluster robust standard errors (Cameron, Miller (2015)) at the three-digit sector level. For each outcome, in column (b) we report the results based on eq. 3.1 and in column (a) those based on a classical two-groups-two-periods Difference in Differences. The outcomes are expressed in logs and include employment, the number of flexible workers and the number of vacancies[13].
Results show that, on average, the positive wage differential paid by non-anticipating firms did not induce them to (differentially) adjust labour costs with changes in the workforce. The coefficient in column (a) indicates a point estimate of the ATT equal to 0.0%. Also the effects on flexible employment (0.4%) and vacancies (0.8%) are non-statistically different from zero. These are usually the drivers of changes in unemployment (Mortensen and Pissarides (1994)), as firms typically stop posting vacancies and do not renew flexible contracts when they plan to cut total employment levels.

[13]Note that we use a proxy for the number of vacancies, given by the inflow of new workers in each firms. For a given firm, a worker is defined as a new inflow in period t if he/she is employed in the firm in period t but wasn't employed in the same firm in period $t-1$. The number of inflows equals the number of vacancies if and only if all vacancies posted by the firms find a match in the labour market. This, despite being a strong assumption in general, is a relatively weaker assumption in crisis periods, when the number of vacancies is typically low and the number of unemployed workers is instead high. The same argument is used in Mortensen and Pissarides (1994) as an argument for their assumption of frictional labour market and the functional form given to the assumed matching function.

Table 3.6: Extensive margin response to nominal rigidity

	dependent variable:					
	Employment		**Flexible Employment**		**N° of Vacancies**	
	(a)	(b)	(a)	(b)	(a)	(b)
started before×reference CAO period	0.000		0.004		0.008	
	(0.010)		(0.017)		(0.029)	
started before×2008Q4		-0.002		-0.006		0.008
		(0.006)		(0.008)		(0.021)
started before×2009Q1		0.022		0.046**		0.010
		(0.014)		(0.021)		(0.025)
started before×2009Q2		-0.002		0.013		-0.033
		(0.009)		(0.021)		(0.030)
started before×2009Q3		-0.010		-0.013		-0.016
		(0.009)		(0.023)		(0.031)
started before×2009Q4		-0.021		-0.029		0.011
		(0.013)		(0.024)		(0.039)
N° of observations (Nt)	1.043.162	1.043.162	1.043.162	1.043.162	1.043.162	1.043.162
N° of firms (N)	90.688	90.688	90.688	90.688	90.688	90.688
N° of CAOs	101	101	101	101	101	101

Note: The table reports the results of the firm-level analysis based on eq. (1). All outcomes are expressed in logarithm. The control variables are: dummy variables for small and big enterprises. the agreed percentage wage increase and the length of the CAO, the sector-level return-on-equity (ROE), return-on-assets (ROA) operating profit margin and value added. Cluster robust standard errors are in brackets. The symbols *, ** and *** denote conventional statistical significance levels.

Table 3.7 reports the result of the analysis of firms' intensive margin responses, obtained again from the estimation of eq. (1). The information on workers' wage bills allows us to obtain a decomposition of workers' wages into all wage components and to distinguish the part attributable to bonuses, benefits and overtime hours[14].

Results show that firms are much more likely to use intensive margin responses to contract staggering rather extensive margin ones: the ATTs from column (a) equal -14.9% for bonuses and -2.3% for benefits. The period-specific ATTs in columns (b) show that the effects on these two wage components turn significantly negative starting from 2009Q2 and 2009Q3, respectively. Instead, no difference in the use of overtime hours worked by the firms is found. Figure 3.3 provides a graphical representation of the effect on employment, bonuses and benefits by plotting the ATTs around the timing of the shock.

[14]Please note that not all the firms in the Netherlands use bonuses, benefits and, especially, overtime hours. For this reason he number of observations differ from Table 3.6, as the log transformation gets rid of all firms that do not pay these compensations.

Table 3.7: Intensive margin response to nominal rigidity

	dependent variable:					
	Bonuses		**Benefits**		**Overtime hours**	
	(a)	(b)	(a)	(b)	(a)	(b)
started before×reference CAO period	-0.149***		-0.023		-0.008	
	(0.057)		(0.018)		(0.033)	
started before×2008Q4		-0.131		-0.017		-0.036
		(0.100)		(0.016)		(0.024)
started before×2009Q1		-0.058		0.010		0.008
		(0.066)		(0.021)		(0.039)
started before×2009Q2		-0.170***		-0.015		-0.060
		(0.078)		(0.019)		(0.039)
started before×2009Q3		-0.161***		-0.064***		-0.072
		(0.060)		(0.017)		(0.047)
started before×2009Q4		-0.266**		-0.043**		0.006
		(0.129)		(0.021)		(0.045)
N° of observations (Nt)	681.663	681.663	649.466	649.466	260.093	260.093
N° of firms (N)	56.805	56.805	54.122	54.122	21.674	21.674
N° of CAOs	101	101	101	101	101	101

Note: The table reports the results of the firm-level analysis based on eq. (1). All outcomes are expressed in logarithm. The control variables are: dummy variables for small and big enterprises. the agreed percentage wage increase and the length of the CAO, the sector-level return-on-equity (ROE), return-on-assets (ROA) operating profit margin and value added. Cluster robust standard errors are in brackets. The symbols *, ** and *** denote conventional statistical significance levels.

Figure 3.3: Extensive and Intensive Margins Response

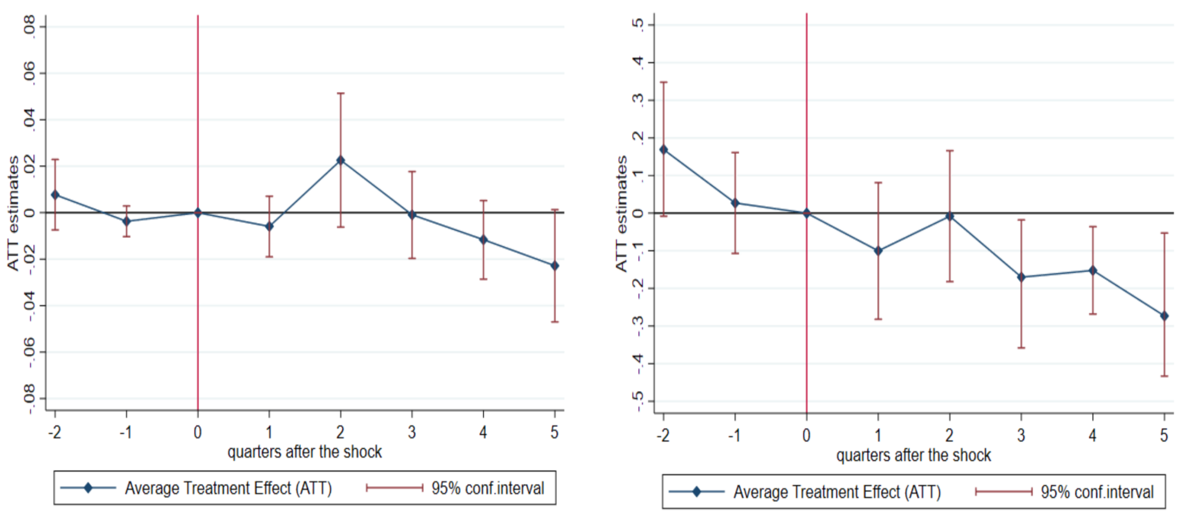

(a) (b)

Note: The figure shows the estimated Average Treatment Effects on the Treated (ATTs) of contract staggering on employment (a) and on bonuses and benefits (b).

Summing up, our results show a precise zero estimate of the effect of contract staggering on employment, and show that that non-anticipating firms are much more likely to compensate the average 2.5 percent higher wage with changes in in the discretionary part of workers' compensation. In sectors where wage resets were agreed before the 2009 recession started, firms did not anticipate the crisis and committed to pay high wage increases, but as soon as the crisis hit they compensated the increase in ordinary wages with cuts in bonuses and benefits, on which single employers have more room for the discretion.

3.4.3 Heterogeneity Analysis

In this section, we perform a heterogeneity analysis to shed more light on the role of contract staggering as a nominal rigidity. The definition of contract staggering not only requires non synchronicity in the validity periods of the various CAOs (that we exploit in our identification strategy) but also requires that the agreed wage increase is valid for a certain number of period (wage stickiness).

Here, we look at the size of the ATTs as a function of contract duration to see whether firms' labour cost adjustments depend on the degree of the wage rigidity. In sectors where collective agreements last longer, firms commit to a certain wage increase for a longer period, and they might have higher incentives to cut labour costs through other margins. On the contrary, firms covered by short lasting CAOs can wait the next bargaining round to directly reduce their wage offer in the CAO. As already shown in Table 3.2, the heterogeneity in contract duration is high in the Netherlands, and varies from six months to about three years.

Table 3.8 reports the result of the heterogeneity analysis, and a graphical representation of the result is given in Figure 3.4. Columns (a)-(c) contain the estimates of the specification in eq. (1) on the subsamples of firms covered by CAOs lasting less than 18 months, between 18 and 30 months, and more than 30 months, respectively. Results show that, as the degree on nominal rigidity increases (i.e. as the contract duration is longer), firms' responses on employment, bonuses and benefits get stronger. We find a negative and significant effect on both employment and bonuses and benefits for firms covered by CAOs lasting more than 30 months. This response is significantly different from the response of firms covered by the most flexible CAOs lasting less than 18 months, as well as from those lasting for a period of 18 to 30 months. At the end of the treatment period, firms' response on the extensive margin is significantly negative, independently from the relative contract duration.

Table 3.8: Heterogeneity Analysis

	dependent variable:					
	Employment			**Bonuses & Benefits**		
	(≤ 18m)	(18-30m)	(≥ 30m)	(≤ 18m)	(18-30m)	(≥ 30m)
started before×2008Q4	0.027*	-0.004	-0.017***	0.012	-0.097	-0.098**
	(0.016)	(0.004)	(0.006)	(0.037)	(0.073)	(0.033)
started before×2009Q1	0.087*	0.015	-0.029*	-0.062	-0.008	-0.167***
	(0.045)	(0.015)	(0.016)	(0.094)	(0.041)	(0.025)
started before×2009Q2	0.037**	-0.005	-0.037***	0.119	-0.173	-0.222***
	(0.017)	(0.013)	(0.011)	(0.207)	(0.105)	(0.027)
started before×2009Q3	-0.017	-0.018	-0.041***	0.068***	-0.129*	-0.192***
	(0.012)	(0.012)	(0.012)	(0.012)	(0.065)	(0.027)
started before×2009Q4	-0.055**	-0.028*	-0.070***	0.089	-0.147	-0.230***
	(0.025)	(0.016)	(0.017)	(0.144)	(0.094)	(0.078)
N° of observations (Nt)	231.307	670.112	141.743	194.191	560.797	106.690
N° of firms (N)	22.239	53.781	14.668	16.051	42.999	9.064
N° of CAOs	31	51	19	31	51	19

Note: The table reports the results of the heterogeneity analysis based on the length of the reference CAOs. All outcomes are expressed in logarithm. For each specification, the first column refers to the subsample with CAOs lasting at most 18 months, the second column to the subsample with CAOs that last more than 18 and less than 30 months. The third column refers to the subsample with CAOs lasting at least 30 months. The firm-level control variables are the same used in the main specification. Cluster robust standard errors are in brackets. The symbols *, ** and *** denote conventional statistical significance levels.

The results of this heterogeneity analysis confirm the presence of relevant rigidities in staggered wage setting. The non synchronicity of the validity periods of collective agreements induces a wage differential across sectors negotiating wage resets around the time of an aggregate shock. We showed that the wage increase paid by firms that did not anticipate the 2009 recession was about 2.5% higher than those of firms that reset their wage during the recession. However, the most relevant rigidity is due to the presence of long contract durations, which imply time commitment of the agreed wage increases. In fact, wages are sticky for the entire validity period of the CAO. We showed that non-anticipating firms covered by CAOs with the longest durations have also the strongest labour cost adjustments. Our main result that contract staggering does not have, on average, real effects on employment contrasts with the results of macroeconomic models such as Gertler and Trigari (2009). In their model calibration for the U.S. economy, they pick a value of the wage reset frequency parameter that implies that firms, on average, reset wages every three quarters. We provide evidence of significant causal effects on employment only in sectors in which wage resets occur less than once every 10 quarters. Our results do suggest that the absence of an overall employment effect is contingent on the use of other adjustment margins at the intensive margin, e.g. bonuses and to a lesser extent benefits.

Figure 3.4: Heterogeneity in Extensive and Intensive Margins Response

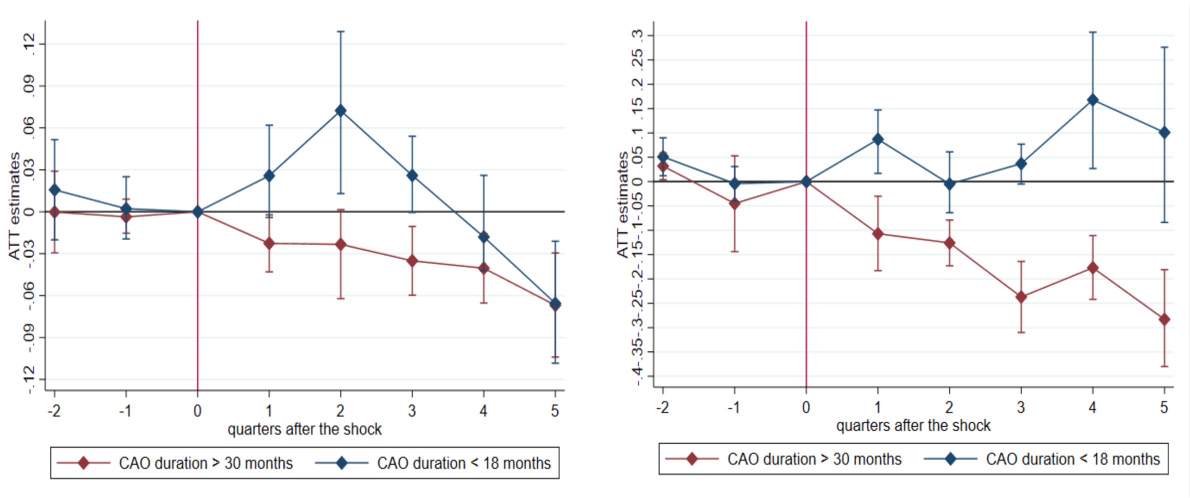

(a) (b)

Note: The figure shows the estimated Average Treatment Effects on the Treated (ATTs) of contract staggering on employment (a) and on bonuses and benefits (b) in sectors covered by CAO having a validity period of less than or equal to 18 months (blue) and more than or equal to 30 months (red).

3.5 Robustness

3.5.1 Agreement vs Start dates

The identification strategy presented in the previous section relies on the fact that CAOs start dates, being pre-determined from the moment the previous contract was signed, are independent on the uncertainty caused by the crisis and thus exogenous.

However, start dates typically differ from agreement dates that are instead endogenous and, during a crisis period, possibly dependent on the uncertainty surrounding it. In the Netherlands, wage resets can be agreed by social partners months after the start date of the collective agreement, and the resulting wage increase can thus apply retroactively (Hijzen et al. 2019). In Figure 3.1 we already showed that indeed the agreement delay has increased during the 2009 recession, as compared to the previous and following period, and this somehow supports the argument of Danziger and Neumann (2005) on the endogeneity of agreement dates.

This issue can potentially affect the causal interpretation of our estimated effects to the extent that firms whose CAO was supposed to start before the crisis, actually agreed the wage increase after the crisis started because of a delay. Once social partners realize that a recession is hitting the economy, they could delay the negotiation process to gather more information or, given the new scenario, even change their wage offer and demand. In this case, considering start dates instead of agreement dates can cause the treatment group to include firms that, because of a delay, managed to anticipate the crisis.

Table 3.9: Agreement dates vs Start dates

N° of collective agreements		Agreed after 2008Q4	Agreed before 2008Q4
Started	after 2008Q4	41	0
	before 2008Q4	10	50

Note: The table reports the number of collective labour agreements started and agreed before or after 2008Q4.

Table 3.9 shows that out of the 60 sectors whose CAO started before 2008Q4, 10 of them reached an agreement after the crisis because of a delay[15]. This means that social partners might have agreed the wage increase in these sectors under a very different information set. To exclude this possibility, in this section we switch from an identification based on start dates to an identification based on both start and agreement dates. The

[15] The average annual wage increase established in this group is equal to 2.35%, and is mostly in line with those started and agreed before the crisis. Out of this 10 CAOs, only two agreed on wage growth close to zero.

crucial difference between the two is that while start dates are pre-determined and thus exogenous, they do not capture the true non-anticipating firms in case of delays in the negotiation process. On the other hand, agreement dates are endogenous but they better reflect the information set under which social parties signed the agreement. Here, we drop the firms covered by the 10 CAOs started before the crisis but agreed during the crisis and we re-estimate eq. (1). In this way, we make sure that the treatment group consists of truly unanticipating firms and, at the same time, we rule out possible issues of selection into treatment.

Table 3.10: Robustness: Identification based on Agreement dates

	dependent variable:					
	Employment		**Bonuses**		**Benefits**	
	(a)	(b)	(a)	(b)	(a)	(b)
agreed before×reference CAO period	0.004		-0.142**		-0.023	
	(0.010)		(0.060)		(0.019)	
agreed before×2008Q4		0.000		-0.114		-0.018
		(0.006)		(0.107)		(0.016)
agreed before×2009Q1		0.024*		-0.055		0.010
		(0.014)		(0.073)		(0.022)
agreed before×2009Q2		0.002		-0.177**		-0.016
		(0.010)		(0.087)		(0.020)
agreed before×2009Q3		-0.011		-0.184***		-0.062***
		(0.010)		(0.066)		(0.018)
agreed before×2009Q4		-0.018		-0.204***		-0.045**
		(0.013)		(0.129)		(0.022)
N° of observations (Nt)	976.294	976.294	642.563	642.563	617.178	617.178
N° of firms (N)	85.050	85.050	51.718	51.718	51.159	51.159
N° of CAOs	91	91	91	91	91	91

Note: The table reports the results of the firm-level analysis based on eq. (1). All outcomes are expressed in logarithm. The control variables are: dummy variables for small and big enterprises. the agreed percentage wage increase and the length of the CAO, the sector-level return-on-equity (ROE), return-on-assets (ROA) operating profit margin and value added. Cluster robust standard errors are in brackets. The symbols *, ** and *** denote conventional statistical significance levels.

Note that, thanks to the availability of information on both the start and agreement dates, the use of the endogenous agreement dates does not bias our result: by excluding the firms that delayed their wage resets, the new treatment group only consists of the firms whose CAO started *and* was agreed prior to October 2008. In such a way, we make sure that the treatment group only consists of non-anticipating firms.

Table 3.10 shows that re-estimating eq. (1) on the basis of the new agreement dates does not modify any of the results obtained in this paper. The point estimates of the effect of contract staggering on both employment and bonus and benefits are exactly equal to the ones obtained in Table 3.6 and 3.7.

3.5.2 Worker-level evidence

The main result of this paper is that firms that did not anticipate the 2009 recession paid a higher contractual wage growth, but then adjusted labour cost via the available adjustment margins, mostly via bonus and, to a lesser extent, via benefits cuts. In this section, we switch to a workers-level analysis to test whether workers employed in treated firms eventually enjoyed a pay increase due to contract staggering. Therefore, we use the employees part of the MEED and we look at the wage bills of all workers that were employed by the 90.688 firms in our sample and that were continuously employed throughout the whole sample period (2007-2009). Monthly wage bills have been collapsed at quarterly frequency too, so that an analogous version of eq. (1) at the workers-level can be obtained. The estimating equation is the following:

$$y_{i,j,s,t} = \lambda_t + c_i + \sum_{k=2008Q4}^{2009Q4} \beta_t (T_s \lambda_k) + \theta \, newCAO_{s,t} + \mathbf{Z}'_{s,t}\gamma + \mathbf{X}'_{i,j,s,t}\delta + u_{i,j,s,t} \tag{3.2}$$

Where $y_{i,j,s,t}$ is the wage of worker i in firm j of sector s at time t, and $\mathbf{X}_{i,j,s,t}$ are workers characteristics. The outcomes we consider are the worker's ordinary and total net hourly wage. The former is equal to the ordinary wage divided by the number of contractual hours, while the latter includes all non-ordinary wage components such as benefits, bonuses and wage from overtime hours worked[16].

Table 11 reports the results of the estimates of eq. (2), for both the cases in which the treatment group indicator T_s is determined using the start dates and the agreement dates of CAOs. Results based on the former dates show that workers employed in non-anticipating firm enjoyed, on average, a 1.4% higher ordinary wage relatively to their control group counterparts, but the difference in the total wage (that includes bonuses

[16]We divide the total net wage in a quarter with the total number of hours worked. In this way, an increase in wage can only be due to an increase in wage, but not to an increase in hours worked.

and benefits) is non-statistically significant. Instead, results based on agreement dates show that after the cut in the discretionary part of the wage bill, the total wage of workers in the treatment group is slightly lower (-0.7%) than that of workers in the control group. We conclude that employees of non-anticipating firms did not eventually earn more as a result of contract staggering and the time of the signature date of their CAO. The fact that individual employers have some room for discretion on the bonus and benefit components gives to the firms a flexibility option to adjust labour cost without changing the workforce.

Table 3.11: Workers' individual wages

	dependent variable:			
	Ordinary hourly wage		**Total hourly wage**	
	(a)	(b)	(a)	(b)
started before × after 2008 Q4	0.014 ***		0.0032	
	(0.003)		(0.0041)	
agreed before × after 2008 Q4		0.008***		-0.007***
		(0.003)		(0.004)
New CAO	0.022 **	0.022 **	0.0150 *	0.0150 *
	(0.009)	(0.009)	(0.0082)	(0.0082)
Age	0.057 ***	0.057 ***	0.0566 ***	0.0566 ***
	(0.001)	(0.001)	(0.0010)	(0.0010)
Age squared	-0.001 ***	-0.001 ***	-0.0005 ***	-0.0005 ***
	(0.000)	(0.000)	(0.0000)	(0.0000)
Experience	0.003 ***	0.003 ***	0.0026 ***	0.0026 ***
	(0.000)	(0.000)	(0.0003)	(0.0003)
Maternity leave	-0.0000	-0.0000	-0.0127 ***	0.0127 ***
	(0.003)	(0.003)	(0.0035)	(0.0035)
Sickness leave	-0.035 ***	-0.035 ***	-0.0393 ***	0.0393 ***
	(0.004)	(0.004)	(0.0037)	(0.0037)
Permanent contract	0.107 ***	0.108 ***	0.1111 ***	0.1111 ***
	(0.008)	(0.008)	(0.008)	(0.008)
N° of observations (Nt)	24.690.219	24.690.219	24.690.219	24.690.219
N° of workers (N)	1.541.430	1.541.430	1.541.430	1.541.430
N° of firms	90.688	90.688	90.688	90.688
N° of CAOs	101	101	101	101

Note: The table reports the results of the worker-level analysis. All outcomes are expressed in logarithm. The control variables are the same used in the main specification. Cluster robust standard errors are in brackets. The symbols *,** and *** denote conventional statistical significance levels.

3.5.3 Parallel Trend test

The assumption behind the results obtained from eq. (3.1) is that, without the nominal rigidities induced by contract staggering, all labour market outcomes in the treatment and control group firms would have had the same trend. In formulas, we are able to consistently identify the Average Treatment effect on the Treated (ATT) only if:

$$E(y_{j,s,a}(0) - y_{j,s,b}(0)|T = 1) = E(y_{j,s,a}(0) - y_{j,s,b}(0)|T = 0) \tag{3.3}$$

Where a and b stand for a generic period after and before the treatment period, respectively. Since the left term is a counter-factual, it cannot be observed in the data, and the condition cannot be tested. This is usually overcome by testing the so-called parallel trend assumption, i.e. by looking at whether the outcomes in the two groups were following similar trends in the period before the treatment.

Figure 3.5: Parallel Trends test

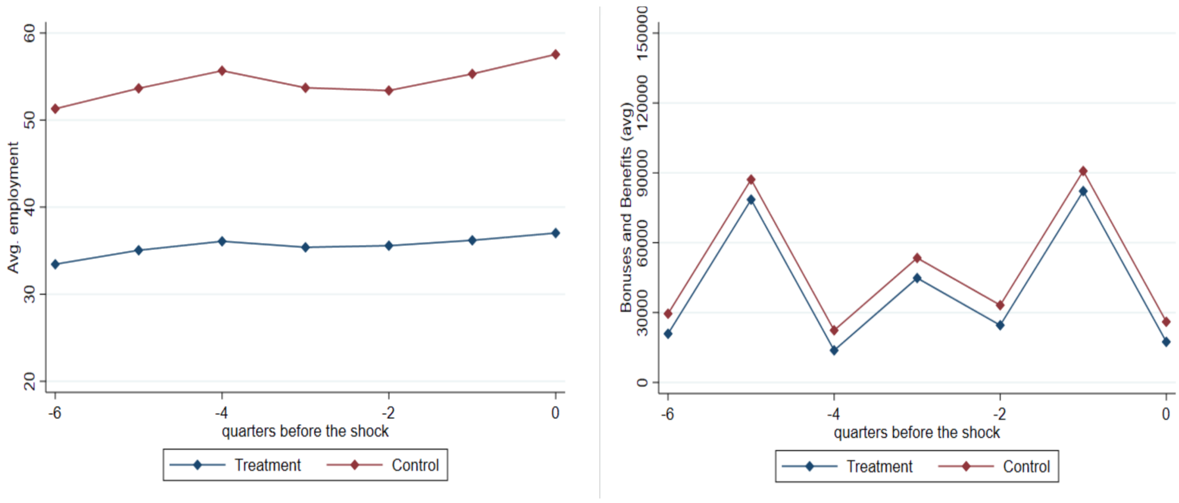

(a) (b)

Note: The figure shows the pre-treatment period trends in the employment level (a) and in the amount paid in bonuses and benefits (b) for firms that signed their CAO before (blue) and after (red) 2018 Q4.

Figure 3.5 reports, for the two groups, the trends in employment and in the amount paid in bonuses and benefits in the seven quarters before the shock (2007Q1-2008Q3). The trends were perfectly parallel in the pre-treatment period, and lead us to conclude that the parallel trend assumption holds and that the ATTs estimated in the previous sections are correctly identified.

3.5.4 Other robustness checks

We perform two additional robustness checks. First, we re-estimate eq. (1) excluding the banking, insurance and financial sector. In fact, firms in these sectors firms typically pay higher bonuses as a share of the corresponding total wages. Also, a possible drop in the bonuses in these sectors may be driven by a drop in stock and financial markets. Second, we check whether results are driven by a specific choice regarding the timing of the shock, and we do this by shifting the treatment period onwards from the 2018Q4-2019Q4 period to the 2019Q1-2020Q1 period. The results, not reported and available upon request, exclude these possibilities and show that all estimates are robust to these changes.

3.6 Concluding remarks

In this paper, we have investigated the consequences of contract staggering in collective labour agreements in the Netherlands. Contract staggering is a labour market rigidity that arises from the combination of wage stickiness and unsynchronized collective labour agreements. We show that the non synchronicity of collective agreements can create considerable wage dispersion across sectors that reset wages in nearby periods, especially in times characterized by widespread uncertainty and/or sudden changes in expectations, such as the 2009 recession. Our descriptive analysis shows that the initial wage growth of CAOs that started before the crisis was about 2.5% higher than that of CAOs starting during the crisis, and shows that other bargaining outcomes such as the average duration and delay of collective agreements have changed too. The increase in the agreement delay is suggestive that the uncertainty surrounding the crisis period might have impacted on wage bargaining, and supports the endogeneity of agreement dates.

In the empirical analysis, we ask whether and how labour market rigidities induce firms to adjust labour cost at the onset of the 2009 recession. To answer this question we exploit the exogenous and staggered start dates of collective agreements to set up a quasi-natural experiment that allows us to identify the causal effect of labour market rigidities.

The main finding of this paper is that, unlike what predicted in macroeconomic models, the pervasive contract staggering in the Dutch labour market did not result in employment losses, even in the aftermath of an unprecedented shock such as the Great Recession. In contrast, we provide evidence that firms adjusted wages on the intensive margin, with bonus and benefit cuts that have fully compensated the higher contractual wage growth of employees of non-anticipating firms. We find that employment cuts have been significant only among firms covered by very rigid collective agreements lasting more than thirty months, and that the cut in bonuses has fully compensated the higher contractual wage growth enjoyed by the employees of non-anticipating firms.

Overall, our results suggest that despite the widespread rigidity, the Dutch wage setting is flexible enough not to cause employment reductions. Hence, our paper underscores the importance of the wider institutional setting in which wage setting takes place. We conclude that the fact that bonuses and benefits are partly left to the discretion of single employers provides a flexibility margin that allows firms to adjust labour cost without employment losses. This is consistent with the evidence from the Wage Dynamic Network initiative (ECB, 2009) showing that bonus and benefit cuts are the most common wage adjustments used by European firms.

Unfortunately, our data does not allow us to point this flexibility to specific aspects of CAOs, e.g. performance pay or profit sharing schemes. Further research is needed to understand which details are more important in enhancing the flexibility of collective agreements. Also, due to data unavailability this paper is silent on other important margins of adjustment that may have played a role. One is certainly firms' pricing: some companies may in fact decide to increase prices rather than adjust labour costs, and this is likely to depend on the level of competition in the sector. Another interesting possibility is the effect on workers' transition to unemployment (in line with Diez-Catalan and Villanueva) and, especially for the Netherlands, workers' transition to self-employment. We leave these questions for future research.

Chapter 4

The Housing Wealth Effect: a comparative analysis of Italy and The Netherlands

4.1 Introduction

We study the housing wealth effect on savings using a comparative empirical approach. Housing represents the largest wealth share in household portfolios in many advanced economies but, unlike other assets, it is hard to decumulate later in the life-cycle. For this reason, while the literature on wealth effect has mainly focused on the estimation of the marginal propensity of consume, the literature on housing wealth effects has extended beyond that in order to highlight the important implications for retirement preparation and retirement decisions.

The topic has recently gained momentum due to the recent set of accommodative monetary policies undertaken by the ECB. Monetary policies affect households decisions via a direct effect (inter-temporal substitution and income effects) and via indirect general equilibrium effects on asset prices. While further income effects may be limited, due the

This chapter is taken from a joint project with Mauro Mastrogiacomo, and the paper has been previously circulating as Caloia, F. G., Mastrogiacomo, M (2020). Pre-pension Savings Out of Housing Wealth: A Comparative analysis of Italy and The Netherlands. Netspar Discussion Paper.

impossibility of further lowering the interest rates at the ZLB, monetary policies can still impact via wealth effects. Especially in some European countries, housing represents the dominant share in household portfolios. Therefore, studying housing wealth effects is key to understand the ultimate consequences of monetary policies on households.

Despite both the macroeconomic and microeconomic literature have paid attention to this relevant topic, the two approaches have always returned different estimates and conclusions. The macroeconomic literature has identified a robust relationship between house prices and consumption at the aggregate level. The microeconomic literature instead finds very mild effects of the same relationship.

Regarding consumption, Attanasio and Weber (1994) study the role of housing in the consumption boom of the UK in the late 1980's and find that housing market explanations account for much of the increase in consumption by the older cohorts, as they are more likely to be homeowners. Also Campbell and Cocco (2007) find a positive effect on consumption for older homeowners and a null effect for young renters. Browning et al (2013) instead find no statistically significant effect. Using time series data for the U.S., Bhatia (1987), Hendershott and Peek (1989) and Skinner (1996) estimate MPCs of around 5%. Regarding savings, Rouwendal and Alessie (2002) and Engelhardt (1996) find a negative association between housing wealth changes and savings in the Netherlands and the U.S., while Suari Andreu (2020) and Disney (2010) find zero or mild effects. The main reason behind the micro/macro differences in the estimate of the marginal propensity to consume out of housing wealth[1] is that, while it is the case that the correlation between consumption growth and house price growth is very high in aggregate data (Case et al. 2005), the micro literature has focused on the causal relation, and thus on excluding competing channels (such as expectations) that could potentially make house price changes endogenous, and thus make the relation with consumption and savings spurious.

In this paper, we adopt a micro approach to also focus on the causal effect of house price changes on household savings decisions. To elicit this causal effect, we decompose the evolution of housing wealth in its exogenous and unanticipated component. Differently

[1] Another reason is that housing wealth is different from financial wealth at least for two reasons: first, housing is an illiquid asset and there are no instruments in the markets (e.g. reverse mortgages) that efficiently allow to liquidate all or part of the home equity; second, housing is also a consumption good that households consume by living in their houses. These factors represent inevitable frictions. Different authors have argued thus that housing wealth could be seen as a sideshow (Skinner, 1996).

from previous studies, we elicit the unanticipated/exogenous part not only filtering out previously stated expectations of house price movements from changes in housing wealth, but also taking into account the role on home-improvements, that represents the endogenous component of housing wealth changes.

To obtain a proper causal estimate of the wealth effect in fact, the change in house value must be not only *unexpected* but also *exogenous*, meaning that it is very important in empirical work to distinguish between changes in prices (exogenous wealth effect) and changes in asset allocations (endogenous wealth effect). For financial wealth, empirical studies have proposed a decomposition of the wealth effect that, for example, takes lagged portfolio weights into account to exclude the portfolio reallocation component from the computation of the wealth effect (see Paiella and Pistaferri, 2017). However, while the lack of consideration of this endogenous component has been mentioned as a bias to micro estimates of the wealth effect (Cooper, 2016), this has not received the same attention in studies that focus on the housing wealth effect. For housing, the endogenous component of the wealth effect can be certainly attributed to home improvements, which can substantially increase the market value of a house, but are inherently unobservable due to the *indivisibility* of houses. We contribute to this literature by showing, theoretically and empirically, the role of maintenance as an additional possible confounder to the estimate of the housing wealth effect. The results of our model show that the role of home-improvements depend, crucially, on the assumption made on house quality. If house qualities are discrete, and reflect key characteristics such as house type or location, then there is little role for home improvements in affecting future home equity. If instead house qualities are continuous, and reflect other characteristics such as energy efficiency, technology, thermal and sound isolation etc. then home improvements are endogenous to the dynamics of home equity. In this case, a shift in the expectations of house prices increases the relative convenience of undertaking home improvements, and ultimately impacts the housing wealth effect. Also, savings and home improvements are shown to be, to some extent, substitutable tools for inter-temporal allocations.

In the empirical part, we adopt a comparative approach to estimate the housing wealth effects in two different countries, using comparable surveys on household income and wealth. The advantage of our approach is to make sure that differences in the estimates are not due to different approach being used, or different information available. The countries that we investigate are Italy and the Netherlands. Despite diametrical differences in housing market characteristics, credit conditions and status of the business cycle, we show how housing wealth effects are actually very similar, but also relatively small. In the empirical part, we show that a 1% unexpected change in housing wealth causes a reduction in active savings in both countries by 0.03% (about 85 euro in the Netherlands and 55 euro in Italy). The estimated magnitude does not change when considering or not home investments. This results show that despite the theoretical relevance of considering home improvements, its empirical relevance is negligible, as the bias generated by omitting this information is very mild. These results suggest that home improvements did not result into a significant increase in the quality of houses, and played no role in the evolution of home equity. The remaining of the study is organized as follows. Section 2 discusses a stylized model with housing and endogenous home improvements. Section 3 discusses the data and methodology. Section 4 and 5 present the descriptive and empirical analysis. Section 6 concludes.

4.2 Theoretical Framework

The aim of this section is to show the role of maintenance and home improvements for the housing wealth effect. The proposed model is a stylized version of Piazzesi and Schneider (2016), who provide a general framework that includes housing in a life-cycle model. Our model preserves at least two key features. First, houses are assets that provide a non-tradable dividend, the housing service, which is a consumption good that allows individuals to derive utility from living in their house. Second, houses are technologies that depreciate if essential maintenance is not performed. Since the main focus of this model (and the empirical section) is on the maintenance, here we do not investigate house purchase decisions.Instead, we focus on homeowners and we endogenize their maintenance decision. Households lifetime utility is represented by the following function:

$$\tilde{U} = U(g(c_t, s(H_t))) + \beta E_t[V(w_{t+1})] \tag{4.1}$$

Where U and V are the strictly concave current and future utility functions, respectively. The function g is an aggregator function that allow distinguishing substitution across goods from substitution across periods. Households derive current utility from consumption and from the housing service, which in turn depends on overall housing quality H_t. A House of quality H_t has a value of $p_t H_t$ and provides a housing service $s(H_t)$. Households derive more utility by living in bigger/higher quality houses, that is $s'(H_t) > 0$. House prices are stochastic and households have expectations over future prices $E_t[p_{t+1}]$. Eventually, In order to show the analytical solution of the model, we will use specific functional forms for the utility functions. For the sake of simplicity, we will assume logarithmic (and thus separable) current utility U and future utility V, and a Cobb-Douglas aggregator g. All analytical derivations are in the appendix.

4.2.1 Exogenous maintenance expenses

As a first step, we assume that house qualities are discrete, that is homeowners are endowed with a house of quality H $\in \mathcal{H}$, where $\mathcal{H}$ is the set of discrete house qualities. As a result of this assumption home improvements, if not extreme, cannot change the stock of quality a homeowner is endowed with. For example, the quality stock of a house reflects characteristics such as location, house type, distance to shops and supermarkets, vicinity to metro and bus stops, presence of good schools in the neighborhood. If this is the case, house quality is endogenous in house purchase decisions, but is exogenous to home improvements decisions. Therefore, we introduce maintenance as periodic expenses aimed at repairing and contrasting the depreciation of houses. Households face the following decision problem:

$$max_{c_t} \quad U(g(c_t, s_t(H))) + \beta_i E_t[V(w_{t+1})] \tag{4.2}$$

$$s.t. \quad c_t + \kappa(H)\varepsilon = y_t - a_t$$

$$w_{t+1} = Ra_t + y_{t+1} + p_{t+1}H$$

The period t budget constraint states that current consumption and maintenance must be paid using labour income y_t, and the remaining endowment is left for savings a_t. The $t+1$ level of wealth is given by future savings, future labour income and future housing wealth, which depends on the realization of house prices, which are stochastic. $\kappa(H)\varepsilon$ are maintenance expenses, represented by a cost that depends on overall house quality $\kappa(H)$ multiplied by a random shock $\varepsilon \sim IID(1, \sigma^2)$. Intuitively, bigger and higher quality houses need more maintenance and are more expensive to maintain, but the size of the expenses depends also on random breaks that occur over time and require fixes. The Euler equation:

$$U'(g(c_t; H_t))g'_c(c_t; H_t) = \beta R\, E_t[V'(w_{t+1})] \tag{4.3}$$

From which, using the assumed functional forms of U,V and g, we obtain the optimal saving rate:

$$a_t^* = \frac{\beta R(y_t - \kappa(H)\varepsilon) - \alpha(y_{t+1} + E_t(p_{t+1})H)}{R(\alpha + \beta)} \tag{4.4}$$

The optimal saving rate depends on the realization of the shock ε. The corresponding housing wealth effect on savings, i.e. the change in savings due to a permanent shift in expectations of future house prices, is given by:

$$\frac{\partial a_t^*}{\partial E_t(p_{t+1})} = -\frac{\alpha}{\alpha + \beta} R^{-1} H \tag{4.5}$$

In words, the change in savings equals a fraction of the present excess house value, where the fraction depends on the discount factor and on the Cobb-Douglas parameter, capturing the marginal utility of current consumption relative to future consumption .

4.2.2 Endogenous home improvements

Next, to introduce endogenous home improvements we assume, instead, that house quality is continuous and additive. In this case, house quality depends also on factors such as energy efficiency and the availability of services such as elevator, car parking, etc. that can be endogenous to home improvement decisions, other than house purchase decisions. Therefore, we assume that houses are technologies that depreciate if essential maintenance is not performed and appreciate if home improvements are undertaken. In particular, homeowners are endowed with a housing stock H_{t-1} and undertake maintenance such that:

$$H_t = H_{t-1}(1-\delta) + h_t \tag{4.6}$$

Where h denotes maintenance undertaken in period t, H_t is total stock of housing quality after maintenance and δ is a depreciation rate. If $h_t = \delta H_{t-1}$ only ordinary maintenance is performed and the overall housing quality stock remains constant. The next period stock of housing quality will be $H_{t+1} = H_t(1-\delta)$.

$$max_{c_t,h_t} \quad U(g(c_t, s_t(H_t))) + \beta_i E_t[V(w_{t+1})] \tag{4.7}$$

$$\begin{aligned} s.t. \quad c_t + \kappa h_t &= y_t - a_t \\ w_{t+1} &= Ra_t + y_{t+1} + p_{t+1}H_t \\ s.t. \quad H_t &= H_{t-1}(1-\delta) + h_t \\ H_{t+1} &= H_t(1-\delta) \end{aligned}$$

In words, households choose consumption and maintenance that maximize lifetime utility. Here, κ is the (unitary) cost of home improvement, relative to the price of numeraire consumption. Since the housing service is increasing in house quality, following Piazzesi and Schneider (2016) we assume that the housing service is equal to its quality, that is $s(H_t) = H_t$. Substituting the budget constraints into the objective function yields:

$$max_{c_t,h_t} U\Bigg(g(y_t - a_t - \kappa h_t; H_{t-1}(1-\delta) + h_t)\Bigg) + \beta_i E_t\Bigg[V\Bigg(Ra_t + y_{t+1} + p_{t+1}\big(H_{t-1}(1-\delta) + h_t\big)(1-\delta)\Bigg)\Bigg] \tag{4.8}$$

The first-order conditions are:

$$\begin{cases} \partial\tilde{U}/\partial a_t: \quad U'(g(c_t; H_t))g'_c(c_t; H_t) = \beta R\, E_t[V'(w_{t+1})] \\ \partial\tilde{U}/\partial h_t: \quad \kappa U'(g(c_t; H_t))g'_c(c_t; H_t) = U(g(c_t; H_t))g'_h(c_t; H_t) + [\beta(1-\delta)E_t(p_{t+1})V'(w_{t+1})] \end{cases} \tag{4.9}$$

The first optimality condition is equivalent to the Euler equation in the case of full certainty, and states that households are indifferent, at the margin, between consuming and borrowing/lending at the risk-free interest rate, as long as $\beta R = 1$. The second optimality condition represents the key optimality condition of this model, and states that households are indifferent, at the margin, between consuming and doing home maintenance as long as the marginal benefit of home improvement is equal to its marginal cost. Taken together, these conditions show that savings and home improvements are substitutable consumption smoothing devices, with two important differences: while savings allow households to reduce current consumption for the sake of increasing future consumption, home improvements allow them to reduce current consumption to increase both current housing services and future home equity. Also, while saving is a risk-free asset, home improvements (and houses in general) are risky assets whose value depends on the evolution of house prices, which follow a stochastic process. However, we do not consider the relative liquidity of these assets, which may differ substantially. Combining the FOCs, we obtain:

$$U(g(c_t; H_t))g'_h(c_t; H_t) = \beta R \left[\kappa - R^{-1}(1-\delta)E_t(p_{t+1})\right]E_t[V'(w_{t+1})] \tag{4.10}$$

This condition governs the home maintenance decision of the household: when deciding whether and how much to invest in home improvements, one trades off additional current housing services and future expected home equity against less current consumption and less future income from the riskless security (the interest rate). Under the assumed utility functions, optimal consumption is equal to:

$$c_t^* = y_t - \kappa h_t^* - a_t^* = \frac{\alpha}{1-\alpha}[R^{-1}(1-\delta)E_t(p_{t+1}) - \kappa](H_{t-1}(1-\delta) + h_t^*) \tag{4.11}$$

Where the optimal maintenance h_t^* is an increasing function of expected future house prices $E_t(p_{t+1})$. Intuitively, when house prices are expected to increase, agents have increasing incentives to make home maintenance, as the investment in home equity is more profitable. Eq. (4.11) shows that, in turn, the effect of a positive house price shock on consumption is positive and proportional to the net present value of the home improvement. Higher house prices induce agents to undertake more maintenance to increase the housing service, and because home equity yields higher expected return relative to savings. Also, increasing house prices increase current consumption via a traditional wealth effect, and thus savings decrease. This implication is also evident from the Euler equation: as home improvements become more profitable, the relative convenience of savings diminishes due to substitutability between savings and home improvements. As a result the housing wealth effect on consumption is given by:

$$\frac{\partial c_t^*}{\partial E_t(p_{t+1})} = \frac{\alpha}{1-\alpha}[R^{-1}(1-\delta)]H_t + \frac{\alpha}{1-\alpha}[R^{-1}(1-\delta)E_t(p_{t+1})]\frac{\partial h_t^*}{\partial E_t(p_{t+1})} \tag{4.12}$$

The first term of eq. (4.12) is the change in consumption attributable to changes in house prices (exogenous component). It is a pure wealth effect, as it equals the expected present value of the increase in home equity. The second term is the part attributable to home improvements (endogenous component) and equals the expected present value of the increase in house quality. Rearranging the budget constraint, we obtain:

$$\frac{\partial a_t^*}{\partial E_t(p_{t+1})} = -\kappa\frac{\partial h_t^*}{\partial E_t(p_{t+1})} - \frac{\partial c_t^*}{\partial E_t(p_{t+1})} \tag{4.13}$$

Which again is negative, as the second and third term are positive. As evident, savings decrease due to the change in consumption, and due to increasing maintenance expenses.

4.2.3 Quality-specific pricing and renovation costs

Let us assume now that house prices are increasing in quality. As an example, energy-efficient houses are more expensive than houses with low energy labels. Similarly, houses in good neighborhoods, or having amenities such as car parking and elevator are also more expensive. Significant home-improvements such as energy saving investments or renovations can also increase the value (price per square meter) of the whole property. We therefore introduce a function $p_t(H_t)$ such that $p_t'(H_t) > 0$. We can write the $t+1$ property value as:

$$p_{t+1}(H_{t+1}(h_t))H_{t+1}(h_t) \tag{4.14}$$

We assume that also renovation costs depend on the size of the improvement, and equal $\kappa(h_t)h_t$. Intuitively, major renovations require obtaining permits, the purchase of materials and the work of a construction company. In this case eq. (4.10), obtained by combining the first-order conditions in (4.9) under assumption (4.14), is given by:

$$U'(g(c_t;H_t))g'_h(c_t;H_t) = \tag{4.15}$$
$$= \beta R \left[\kappa'(h_t) - R^{-1}(1-\delta)\left(p'_{t+1}(H_{t+1}(h_t))H_{t+1} + p_{t+1}(H_{t+1}(h_t)\right)\right] E_t[V'(w_{t+1})]$$

As $H'_{t+1}(h_t) = 1-\delta$ and $H'_t(h_t) = 1$. This last equation is conceptually identical to eq (4.10), and states that households are indifferent between saving and investing in home improvements, if the cost of home improvement is equal to its present expected value. However, the notable difference is that, under the assumption of quality-specific house prices, investing in home maintenance is generally more convenient, in fact investing in home-improvements (i) increases the quality stock of the house and (ii) augments the value of every unit of housing (quality).

The empirical implications we draw are that, if home improvements are endogenous, (i) expectations of increasing house prices affect the propensity for home investments. If these are undertaken,(ii) maintenance expenses have a direct impact on savings, via a simple budget effect and (iii) an indirect impact on savings via the wealth effect, whereby home investments endogenously increase the quality and the value of houses. Therefore, to estimate the MPC/MPS out of housing wealth, it is important to disentangle the effect of home improvements.

4.3 Data and Methodology

For a description of Italian and Dutch household portfolios, we use the Household Finance and Consumption Survey (HFCS) data. The HFCS is an initiative coordinated by the ECB, where national central banks gather micro data using the same survey questions. The survey is based on 84,000 interviews conducted in 18 euro-area countries. In the empirical analysis, we use the two surveys on which the HCSF is based. These describe the same populations of the HFCS, but contain several additional questions which are not present in the HFCS. These questions are very similar in both surveys, and thus allow testing our model without concerns that results will be driven by methodological differences.

For the Netherlands, we use the DNB Household Survey (DHS). The DHS is an annual survey representative of the Dutch speaking population and contains information on income and wealth, as well as on all the psychological aspects of financial behavior. The data is administered by Center Data, on behalf of the Dutch National Bank (DNB). In the DHS, data on expectations of future house prices are available from 2003 for homeowners, while questions about maintenance and home improvements have been asked to the respondents since 2012. In the empirical part that follows, we will concentrate on the 2012-2018 waves. Comparison with the Italian data is though only possible in the period until 2014 as these surveys contains similar questions, with particular reference to those on expectations and home improvements. In the DHS, the expectation question is straightforward and asks *"How much percentage points a year will they increase/decrease on average?"*

For Italy, we use the Survey on Household Income and Wealth (SHIW) administrated by Bank of Italy, which is a biannual survey also containing detailed information on income and wealth of a representative sample of Italian households. In the SHIW, respondents of the 2010 and 2012 wave are asked about their expectations on future house prices, while the question on extraordinary house maintenance has always been present in the questionnaire. Regarding the question on expectations, in the SHIW there are a number of issues that we have to address. The first issue is that respondents of the 2010 questionnaire of the SHIW are not directly asked to report their own subjective expectations, but instead they have been asked the following questions:

- "*On a scale from 0 to 100, what is the probability that house prices will drop in the next 12 months?*"
- "*And what is the probability they will drop by more than 10%?*".

Let Ω_t be the agent's information set at time t. Then, following Paiella and Pistaferri (2017) and assuming that expectations follow a standard normal distribution, it is possible to retrieve the first two moments (mean and variance) of the distribution of expectations by solving the following system of equations:

$$\begin{cases} pr(r_{t+1} < A|\Omega_t) = \Phi\left(\frac{E_t(r_{t+1})-\mu}{\sigma}\right) \\ pr(r_{t+1} < B|\Omega_t) = \Phi\left(\frac{E_t(r_{t+1})-\mu}{\sigma}\right) \end{cases} \tag{4.16}$$

Where Φ denotes the Cumulative Density Function (CDF) of the Standard Normal Distribution, and A and B are the level of returns mentioned in the expectation questions, $pr(r_{t+1} < A|\Omega_t)$ and $pr(r_{t+1} < B|\Omega_t)$ are the observed data points. The unknowns of this system of equations are the expectations $E_t(r_{t+1})$ and the standard deviation σ.

A second issue is that in the 2012 wave, by means of randomization, the same question was asked to a 50% subsample of respondents, while to the other 50% subsample was asked to distribute 100 points among the possibilities that in the next 12 month house prices will be (i) much higher (more than 10%), (ii) slightly higher (between 2% and 10%), (iii) about the same (between -2% and 2%), (iv) slightly lower (between -2% and -10%) or (v) much lower than today (less than -10%). In this case, we obtain the corresponding expectation by assigning the elicited probability weights to the midpoints of each answer category, that is:

$$E_t(r_{t+1}) = \sum_k p_k r_k \tag{4.17}$$

A third issue is that the SHIW data is a biannual survey but the expectation question asks respondent to elicit their one-year-ahead beliefs. To overcome this, Paiella and Pistaferri (2017) show that it is possible to rewrite the expected change in wealth as:

$$E_t(W_{t+2}) = W_t E_t(1 + r_{t,t+1})(1 + r_{t+1,t+2}) \tag{4.18}$$

Due to the observational gap $E_t(r_{t,t+1})$ is observed while $E_t(r_{t+1,t+2})$ is not. Assuming that individuals have AR(1) expectations, then $E_t(r_{t+1,t+2}) = \rho E_t(r_{t,t+1})$ where ρ is the autoregressive parameter. Under this assumption the change in wealth becomes:

$$E_t(W_{t+2}) = W_t(1 + E_t(r_{t,t+1}))(1 + \rho E_t(r_{t,t+1})) \tag{4.19}$$

Suari Andreu (2020) and Brauning et al (2013) show that self-reported expectations closely match data generated via an AR(1) process. However, since the time series dimension of survey data is typically not long enough to fit a model to retrieve an estimate of u, we follow Paiella and Pistaferri (2017) and we assume the autoregressive parameter to be equal to one[2], given the higher persistence of house prices relative to financial assets.

[2]They also show that results are unchanged when other values for the autoregressive parameter are assumed.

A final issue is that, unlike for financial wealth, it is not possible to decompose the housing wealth effect in the endogenous and exogenous component. In fact, the change in the value of a portfolio composed by J different assets can be written as:

$$\begin{aligned}\Delta W_{t,t+1} &= \sum_j W^j_{t+1} - \sum_j W^j_t \\ &= \sum_j p^j_{t+1}(A^j_{t+1} - A^j_t) + \sum_j (p^j_{t+1} - p^j_t)A^j_t\end{aligned}$$

So, it is possible to decompose the change in the value of the portfolio due to a change in the value of the constituent assets and due to a portfolio reallocation. This equation can be easily estimated as price changes and changes in allocation can be observed. Similarly, using the same notation used in Section 2, we can write the change in housing wealth as:

$$\Delta W_{t,t+1} = p_{t+1}H_{t+1} - p_t H_t = (p_{t+1} - p_t)H_t + p_{t+1}h_{t+1} \tag{4.20}$$

Where $p_{t+1}h_{t+1}$ is the value of the home improvement of quality h_{t+1}, which is inherently unobservable due to the indivisibility of the quality and value of houses.

4.4 Descriptive Analaysis

When comparing the portfolios of Dutch and Italian households, large differences emerge. Table 4.1 reports the average portfolio shares, as well as the mean and median values of assets and liabilities, as reported in the HFCS. The table shows important differences and similarities between Italy and the Netherlands: while home-ownership and the value of the main residence are comparable across countries, financial assets holdings greatly differ. Italian households less often have a saving account, and when they do, they have less saving into it. The same is also true for mutual funds and private pensions. At the same time, the share of Dutch households with mortgage debt is almost 40% higher in the Netherlands and, conditional on having a mortgage, the outstanding debt amount tends also to be much higher in the Netherlands.

Table 4.1: Descriptive statistics

	dependent variable:					
	Ownership (%)		**Mean values**		**Median values**	
	(IT)	(NL)	(IT)	(NL)	(IT)	(NL)
Main residence	71%	74%	244 970	296 205	200 000	255 000
other real estate properties	3%	0%	158 765	240 639	99 000	127 500
Business wealth	18%	5%	160 557	152 041	20 000	80 000
Deposits	26%	87%	14 452	34 523	6 808	14 676
Mutual funds	7%	23%	40 676	33 569	20 000	10 724
Bonds	17%	8%	43 735	38 728	20 627	17 000
Voluntary pension/life insurances	14%	51%	14 560	48 884	10 000	19 474
Liabilities	23%	68%	40 892	141 789	13 000	107 000
Mortgage debt	9%	55%	71 923	154 337	50 000	127 000
Other debts	17%	31%	16 206	36 768	5 600	12 669
Credit lines/overdraft	3%	18%	5 275	6 807	2 500	2 000

Note: Descriptive statistics of Italian and Dutch households' portfolios. Mean and median values are conditional on ownership. Source: HFCS.

Housing market developments in Italy and the Netherlands have widely differed in the past decade. Figure 1 shows that house price growth rates in Italy always stayed behind those of the Netherlands, where price dropped after 2010-2011 but recovered after 2013.

Figure 4.1: House price trends in Italy and the Netherlands

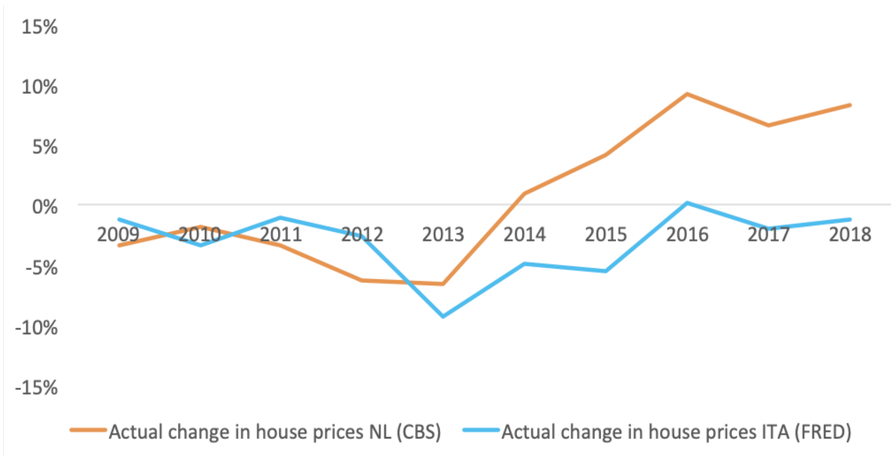

The drop has also been stronger in Italy, where house price growth rates reached about -10%. In the Netherlands, the drop has been about half the Italian one, and was close to -5%. On the contrary, the recovery has been stronger in the Netherlands where, in 2016, the house price growth rate has been close to +10%. Indeed, the housing market in the Netherlands has always been more volatile than the housing market in Italy.

Figure 4.2: Expected vs realized house price changes

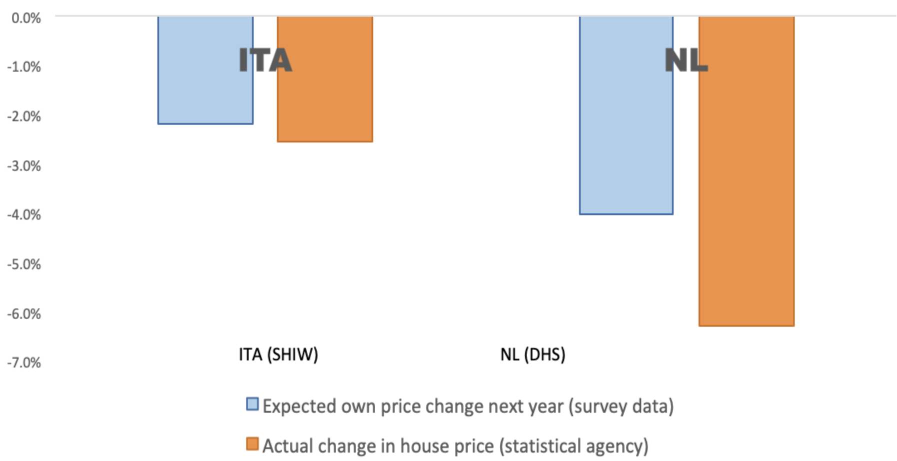

In Figure 4.2, we propose a descriptive comparison of the subjective expectations of changes in house prices. It shows that, on average, Italian respondents reported expectations have matched the actual price change quite closely: the expected price change was equal to -2.2%, close to the realized price change (-2.6%). In the Netherlands, respondents have underestimated the change in house prices: the average expectation was -4.0% while the (average and self-reported) change in the house value over the past two years (i.e. during the 2010-2012 period) was -6.3%. This does not necessarily mean that Dutch respondents are less precise, but could simply reflect the higher volatility of the Dutch housing market.

Figure 4.3: Distribution of expected house price changes

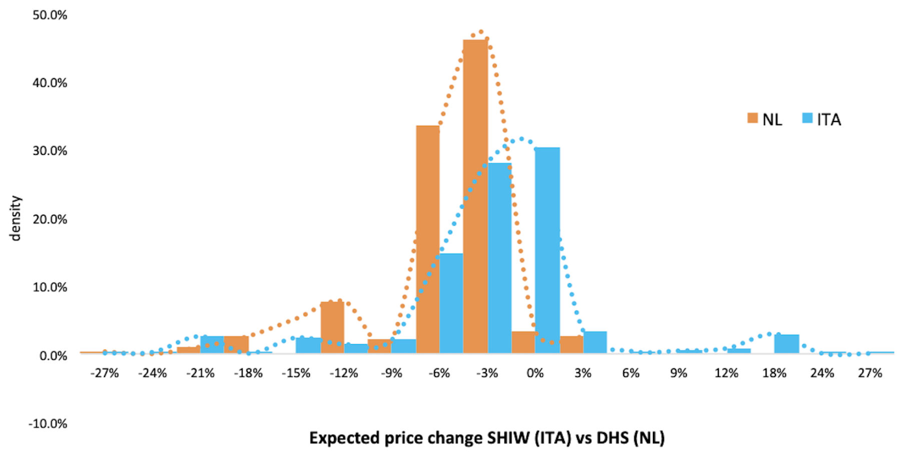

Concerning the distribution of expectations, Figure 4.3 shows that, although both expectations were negative, in 2012 Italian respondents were slightly more positive than the Dutch ones with respect to future house price developments: the average expectation is about -1% in Italy and -3.5% in the Netherlands. The figure also shows that the expectations of Italian respondents display much more variability than those of Dutch respondents. The figure shows that the house price expectations of the Dutch are more left-skewed than those of Italians, which again confirms the relative more pessimistic view of the Dutch in the 2012 survey.

Figure 4.4: Maintenance expenses

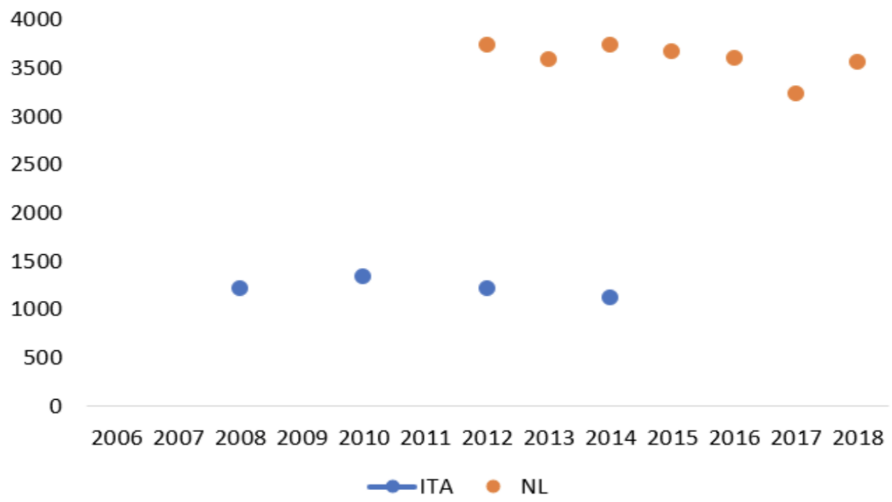

Figure 4.4 shows the mean costs of home improvements for Dutch and Italian households. It shows that Dutch households spend on average twice as much on home improvements, conditional on having spent for maintenance. However, the share of households that do make home improvements differs greatly between the two surveys. In Italy, 20% of homeowners report maintenance, while in the Netherlands this share is about 50%. This difference could be due to the fact that Dutch households who live in a condominium pay a periodical maintenance fee, in order to save for future maintenances.

4.5 Empirical Analysis

The aim of this section is to empirically investigate the role of maintenance in the estimation of the housing wealth effect. Our starting point is that home improvements change the quality of the house and, thus, endogenously affect the value of the property. However, as shown in eq. (16), disentangling the exogenous and endogenous component for housing wealth is not possible (unlike for financial wealth). The reason is that, because of the indivisibility of houses, the term $p_{t+1}h_t$ is unobservable. Instead, we observe the cost of maintenance κh_t, according to the notation of section 2. The literature has mostly relied on the following estimating equation:

$$\Delta S_{i,t} = \alpha + \beta_u\, log(\Delta W_{i,t} - E_{t-1}(\Delta W_{i,t})) + \beta_e\, log(E_{t-1}(\Delta W_{i,t})) + \mathbf{X}'_{i,t}\gamma + \epsilon_{i,t} \qquad (4.21)$$

Where $S_{i,t}$ represents active savings of household i in wave t, which is elicited by asking respondents how much money has been put aside in the last 12 months. Active savings is a measure of savings that is not attributable to capital gains. The two right-hand side variables are the unexpected and expected change in housing wealth, respectively. The expected change in wealth is computed as in eq. (4.20), while the realized wealth change is computed as the difference in the self-reported house values between two waves. $\mathbf{X}'_{i,t}$ are control variables. The advantage of specification (4.21) is that, by filtering out the expected change in housing wealth from realized changes, it allows to estimate the housing wealth effect. According to the life-cycle model, only unexpected changes in the value of wealth should translate into corresponding changes in consumption or savings.

However, part of the realized change in wealth can still result from undertaken home improvements (endogenous wealth effect), other than changes in house prices (exogenous wealth effect). To investigate the role of maintenance, we treat the previous estimation as an omitted-variable problem, and we augment it to account for home improvements:

$$\Delta S_{i,t} = \alpha + \beta_u log(\Delta W_{i,t} - E_{t-1}(\Delta W_{i,t})) + \beta_e log(E_{t-1}(\Delta W_{i,t})) + \delta m_{i,t-1} + \mathbf{X}'_{i,t}\gamma + \epsilon_{i,t} \quad (4.22)$$

Where $m_{i,t-1}$ denotes maintenance undertaken between waves $t-1$ and t. Adding maintenance as a conditioning variable is consistent with our model: maintenance expenses certainly affect savings directly via the budget constraint, also intertemporally (installment payments are common), and possibly they influence the change in house values indirectly via quality accumulation. Due to the indivisibility of house quality and value, house price changes may be correlated with the size of the maintenance expense. We use the cost of home-improvements κh_t as a proxy of the value of home-improvement $p_{t+1}h_t$ which, following eq. (4.20), would allow to disentangle the endogenous component of the wealth change. If measurement error is of classical form, this would induce an attenuation bias in the coefficient δ, capturing the direct effect on savings of additional maintenance costs. However, the object of interest is the housing wealth effect: if the true population model is given by eq. (4.22), estimating eq. (4.21) without taking into account maintenance can possibly lead to a bias in the estimated housing wealth effect. The bias is:

$$E(\widehat{\beta}_u) - \beta_u = (\Delta W_{i,t}^{u'} \Delta W_{i,t}^{u})^{-1} E\big[\Delta W_{i,t}^{u'}\, m_{i,t-1}\big]\delta \tag{4.23}$$

According to our model, the covariance is supposed to be either positive or zero. If home improvements translate into an effective increase in the quality of houses, than $\widehat{\beta}_u$ would be overestimated. Instead, if maintenance activities do not contribute to increase the overall quality of houses, then maintenance costs and realized wealth changes are independent, and the parameter is consistently estimated.

Here, we present estimates of the housing wealth effect on savings based on eq. (4.21) and (4.22) to shed more light on this mechanism. In particular, we study whether estimates of the housing wealth effect that neglect the role of maintenance may return results which are substantially different than estimates that do consider home improvement. Tables 4.2 and 4.3 report the results of the empirical analysis estimated on both the SHIW and the DHS data. In specification (a) we estimate the traditional equation for the housing wealth effect, as in Paiella and Pistaferri (2017). In specification (b) we account for the cost of home improvement. Specification (c) is analogous to (b), but uses a proxy of the value of the home improvement, instead of the cost. This is obtained by compounding the cost of maintenance with the expected change in house prices. In specifications (d) and (e) we re-estimate the specifications (b) and (c) by instrumenting the cost and the value of maintenance. In fact, if measurement is of classical form, as it seems reasonable to assume (see also Juster et al, 2005), this would result in an attenuation bias in the associated coefficient, capturing the direct effect of maintenance on savings. If however, measurement error is correlated with the change in house value, it would affect the estimation of the housing wealth effect.

For Italy, we take advantage of a a tax benefit that allows deducting 55% of the expenses for energy saving investments, and 36% of the maintenance expenses, and we instrument maintenance expenses with the amount of tax deducible expenses. The instrument exploits difference in the deduction rates, and in the maximum deductible amount across investment types, as well as in the application for tax deduction across households[3]. For the Netherlands, we use a dummy equal to one if the household head states that the home improvement "*will reflect fully in an increase of the property value in the case of a sale*". While the two instruments are different in nature, they are conceptually similar: lower maintenance costs (due to tax deduction), as well as higher expected return out of the investment (due to expectations of increasing home equity) should increase the propensity of undertaking maintenance. Therefore, they should correlate with maintenance expenses, while being uncorrelated with the error term, and only indirectly affecting savings decisions.

There are three main results evident from Tables 4.2 and 4.3. The first result is that, despite the large differences documented in section 4.4, similar responses to unexpected changes in wealth emerge in the two countries. Most importantly, these are relatively low in magnitude: a 1% unexpected increase in housing wealth causes a 0.03% reduction in active savings for both countries. Our back of the envelope calculations show that for the Netherlands (Italy) a 2800 (1800) euro unexpected increase in housing wealth causes a reduction in active savings of about 85 (55) euro. Equivalently, the exponentiated coefficient indicates a one euro unexpected change in housing wealth leads to a decrease in savings by three cents. Our result is somehow in between the macro and micro literature for the size of the estimated housing wealth effect. The micro literature in fact finds no or very mild effect, with MPCs in the range of 0%-3% (Attanasio et al. (2009), Disney et al. (2010), Browning et al. (2013)), while the macro literature typically finds larger MPCs out of housing wealth ranging from 5% to 7% (Campbell and Cocco (2007), Case

[3]The law establish a 55% deduction of energy saving investments up to 100.000 euro, 55% up to 60.000 for thermal insulation investments, and a 55% up to 30.000 for heating systems and a 36% deduction for home improvements up to 48.000 euro. Figures refer to 2012.

et al. (2007), Bathia (1987), Skinner (1996)).
Moreover, this result is in line with the evidence that housing wealth effects tend to be smaller than financial wealth effects, due to the fact that (i) there are few instruments on the market that allow households to consume out of home equity and (ii) houses are consumption goods, other than assets, which make households less sensible to fluctuations in housing wealth than to fluctuations in financial wealth.

Table 4.2: Results: the Netherlands

	Dep. variable: Change in Active Savings				
	(a)	(b)	(c)	(d)	(e)
Expected change in housing wealth	0.00913	0.00882	0.00884	-0.0113	-0.0117
	(0.0104)	(0.0107)	(0.0107)	(0.0280)	(0.0218)
Unexpected change in housing wealth	-0.0276***	-0.0275***	-0.0275***	-0.0242**	-0.0241**
	(0.00910)	(0.00911)	(0.00911)	(0.0101)	(0.0101)
Value of maintenance		0.00329		0.2255	
		(0.0257)		(0.2897)	
Cost of maintenance			0.0330		0.2252
			(0.0257)		(0.2895)
Age head 18-34	0.216	0.221	0.213	0.975	0.976
Age head 45-44	0.312	0.321	0.311	0.778	0.783
Age head 45-54	0.022	0.026	0.021	0.064	0.071
Age head 55-64	0.034	0.030	0.033	-0.023	-0.011
Married	-0.070	-0.083	-0.084	-0.040	-0.043
Divorced	-0.157	-0.160	-0.166	0.246	0.202
Head employee	0.092	0.090	0.095	0.574	0.584
Head unemployed	-0.039	-0.041	-0.038	0.601	0.635
Head low education	0.029	0.025	0.029	0.201	0.193
Head high education	-0.109	-0.106	-0.108	-0.385	-0.391
Male head	0.013	0.015	0.017	-0.445	-0.438
Household disposable income	-0.037*	-0.037*	-0.037*	-0.017	-0.017
Net financial wealth	0.013	0.013	0.013	0.056**	0.056**
Constant	-0.084	-0.098	-0.097	-1.119	-1.113
Estimate	OLS	OLS	OLS	2SLS	2SLS
First stage F stat.	-	-	-	28.9	28.9
Observations	3,605	3,598	3,605	1,527	1,528

Note: The top panel of the table reports the coefficients associated to the housing wealth effect. The bottom panel reports the coefficient associated to the control variables. Standard errors in brackets. The symbols *, **, and *** denote ten, five and one percent statistical significance levels, respectively.

Table 4.3: Results: Italy

	Dep. variable: Change in Active Savings				
	(a)	(b)	(c)	(d)	(e)
Expected change in housing wealth	-0.0496***	-0.0502***	-0.0503***	-0.0521***	-0.0515***
	(0.0155)	(0.0155)	(0.0155)	(0.0175)	(0.0175)
Unexpected change in housing wealth	-0.0340***	-0.0342***	-0.0342***	-0.0347***	-0.0348***
	(0.0129)	(0.0129)	(0.0129)	(0.0135)	(0.0135)
Value of maintenance		-0.0759**		-0.2403	
		(0.0363)		(0.0398)	
Cost of maintenance			-0.0741**		-0.2433
			(0.0362)		(0.0396)
Age head 18-34	-0.417	-0.664	-0.495	-0.538	-0.376
Age head 45-44	-0.594	-0.739	-0.620	-0.525	-0.427
Age head 45-54	-0.187	-0.227	-0.217	-0.294	-0.206
Age head 55-64	-0.108	-0.288	-0.139	0.074	0.197
Married	-0.111	-0.091	-0.091	-0.109	-0.117
Divorced	0.082	-0.062	0.101	-0.069	0.050
Head employee	0.229	0.236	0.234	0.373	0.354
Head unemployed	-0.017	-0.040	-0.014	0.000	-0.012
Head low education	0.223	0.179	0.187	0.238	0.235
Head high education	1.054**	1.166***	1.111***	1.045**	0.906**
Male head	0.037	0.012	0.056	0.009	0.093
Household disposable income	-0.069	-0.028	-0.044	-0.016	-0.037
Net financial wealth	-0.049***	-0.049***	-0.048***	-0.051***	-0.051***
Constant	0.857	0.598	0.683	0.626	0.760
Estimate	OLS	OLS	OLS	2SLS	2SLS
First stage F stat.	-	-	-	82.3	84.5
Observations	2,741	2,741	2,741	2,741	2,741

Note: Standard errors in brackets. The symbols *, **, and *** denote ten, five and one percent statistical significance levels, respectively.

The second result is that, when adding the value or the cost of home improvements in specifications (b) and (c), the coefficient associated to these variables are negative and significant in Italy, and zero in the Netherlands. That is, only among Italian respondents there is evidence of a direct negative impact on savings of increasing maintenance costs. Moreover, in both countries the pure housing wealth effect (i.e. the unanticipated component) does not change in magnitude with respect to specification (a). This suggests that home improvements did not translate into significant increase in the quality and the value of houses, and that the size of the bias in eq. (4.23) is very mild or null. According to the different versions of our model, there are three possible explainations. The first reason is maintenance activities are exogenous to the evolution of home equity. The second is that in this case (but not in general), maintenance activities were aimed only at limiting depreciation. The third is that, for major renovations, the investment cost is larger than the present value of the increase in home equity. In all these cases, there is no indirect effect on the quality and the value of houses, and thus there is no endogenous component in the housing wealth effect. Figure 4.5 shows that indeed the correlation between the cost of home improvement and the change in housing wealth is positive (as expected) but very mild.

Figure 4.5: Correlation between size of home improvements and changes in house values

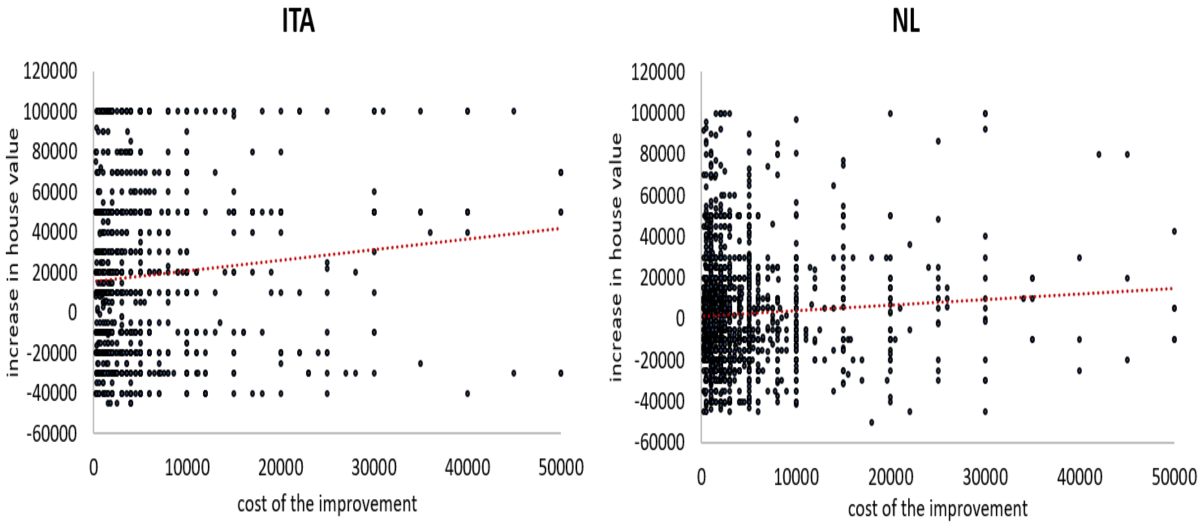

The third result is that, when instrumenting the cost or the value of maintenance, the estimated housing wealth effect does not change in magnitude, while the point estimate of the coefficients becomes many times larger. This evidence suggests that there seems to be no correlation, due to measurement error of maintenance costs, between changes in the value of houses and the error term. Alternatively, it suggests that measurement error is likely to be of classical form. Eventually, the last result is that in Italy, we find a negative and significant effect on savings of expected wealth changes. This result, despite contradicting the life-cycle model that predicts that only unexpected changes should lead to a revision of consumption and savings decisions, is in line with the findings of Paiella et al (2017) of a positive effect on consumption on the same SHIW data. They justify this result appealing to credit constraints.

4.5.1 Heterogeneity Analysis

In this section, we perform a heterogeneity analysis to see how the estimated wealth effects vary in various subgroups of respondents in the two countries. Given the limited amount of observations available, most subgroups are formed on the basis of median values of the variables of interest. The estimates are reported in table 4.4 and are the result of specification (c).

Table 4.4: Heterogeneity analysis

	Unexpected Change in Housing Wealth			
	Italy		**The Netherlands**	
	coef.	*S.E.*	*coef.*	*S.E.*
All (baseline)	-0.0342***	0.0129	-0.0275***	0.0091
No home improvement	-0.0379***	0.0136	-0.0113	0.0155
Home improvement	-0.0242	0.0381	-0.0353***	0.0114
Income below median	-0.00757	0.0160	-0.0171	0.0110
Income above median	-0.0672***	0.0208	-0.0373**	0.0146
Financial wealth below median	-0.0302*	0.0160	-0.0225	0.0145
Financial wealth above median	-0.0500**	0.0204	-0.0304***	0.0112
Age below median	-0.0481**	0.0188	-0.0318**	0.0147
Age above median	-0.0235	0.0179	-0.0234**	0.0109

Note: The table reports the coefficient and the standard error of the estimate housing wealth effect, using specification (c) of Table 2 and 3. The symbols *, ** and *** denote conventional statistical significance levels.

The results show substantial heterogeneity across groups. In both countries, the highest estimated wealth effects are on those with above median income and wealth. Younger respondents also display larger wealth effects. The main cross-country difference emerges when analyzing the role of maintenance, by comparing the estimated effects on those who make home improvements or not. In the Netherlands instead, the housing wealth effect is significant only among respondents undertaking home improvements, and adding maintenance do not attenuates the coefficient. In Italy, the housing wealth effect is driven by respondents that don't do home improvements, while in the other group the coefficient is non-significantly different from zero, mainly because standard errors widen due to the fall in the number of observations (20% of Italian respondents report home improvements).

4.6 Conclusions

This paper sheds light on the role of home improvements in the estimation of the housing wealth effect, by means of a comparative study of Italy and the Netherlands. Our contribution is to show an endogeneity channel in the estimation of the housing wealth. We first propose a theoretical model in which we endogenize the home improvement decision of the household. The model yields a substitutability between savings and home improvements for inter-temporal allocations, and shows that part of the housing wealth effect is endogenous and attributable to home improvements. In fact, expectations of increasing house prices not only induce households to consume more and save less via a traditional wealth effect, but induce agents to undertake more home improvements, as housing investments are relatively more convenient than savings in periods of increasing house prices. The endogeneity channel is exactly due to the fact that, by undertaking more maintenance, agents directly affect housing wealth changes via an increase in house quality. In the empirical part, we show how not considering home improvements decision can affect the estimate of the housing wealth effect. The value added of our comparative approach is that we apply the same methodology to comparable data of two countries, which differ substantially in the characteristics of their housing markets. Our approach also provides an estimate of the housing wealth effect that reflects the exogenous and unanticipated component of the wealth change, obtained by decomposing the unanticipated and anticipated component, and accounting for maintenance costs. We first show that the size of the bias is proportional to the covariance between expenses in home improvements and changes in wealth. Then, we provide evidence that the size of this bias is quantitatively small in the data, because of the mild correlation between the two. The overall results of the paper suggest that home improvements did not translate into significant increases in the quality and the value of houses,

Appendix: analytical derivations

Exogenous Maintenance: Under the assumption of logarithmic utilities U and V, and assuming a Cobb-Douglas aggregator g, the objective function in eq. (4.1) becomes:

$$\begin{aligned}\tilde{U} &= U(g(c_t, s(H))) + \beta E_t[V(w_{t+1})] \\ &= log(c_t^{\alpha} s(H_t)^{1-\alpha}) + \beta E_t[log(w_{t+1})] \\ &= K + \alpha log(c_t) + \beta E_t[log(w_{t+1})]\end{aligned}$$

Where K is a constant. Maximizing this objective function subject to the constraints in eq. (4.2) yields the following first-order condition:

$$\frac{\partial \tilde{U}}{\partial a_t} : \quad \frac{\alpha}{c_t} = \beta R \frac{1}{w_{t+1}}$$

Substituting the budget constraints and rearranging yields the saving rate in (4.4).

Endogenous home improvements: Under the assumption of logarithmic utilities U and V, and assuming a Cobb-Douglas aggregator g, the objective function in eq. (4.1) becomes:

$$\begin{aligned}\tilde{U} &= U(g(c_t, s(H_t))) + \beta E_t[V(w_{t+1})] \\ &= log(c_t^{\alpha} s(H_t)^{1-\alpha}) + \beta E_t[log(w_{t+1})] \\ &= \alpha log(c_t) + (1-\alpha) log(s(H_t)) + \beta E_t[log(w_{t+1})]\end{aligned}$$

Maximizing this objective function subject to the four constraints in eq. (4.7) yields the following first-order conditions:

$$FOC_1 = \frac{\partial \tilde{U}}{\partial a_t} : \quad \frac{\alpha}{c_t} = \beta R \frac{1}{w_{t+1}}$$

$$FOC_2 : \frac{\partial \tilde{U}}{\partial h_t} : \quad \kappa \frac{\alpha}{c_t} = \frac{1-\alpha}{H_t} + \beta E_t(p_{t+1})(1-\delta) \frac{1}{w_{t+1}}$$

Where c_t and w_{t+1} are defined in the t and $t+1$ budget constraints, respectively. Plugging FOC_1 into FOC_2 yields the optimal saving rate a_t^*, as a function of the optimal investment in maintenance:

$$a_t^* = y_t - \kappa h^* + \frac{\alpha}{1-\alpha}[R^{-1}(1-\delta)E_t(p_{t+1}) - \kappa](H_{t-1}(1-\delta) + h_t^*)$$

Using the equation above, it is possible to obtain eq. (4.11) and (4.12) via the period t budget constraint. Next, plugging this last equation back into FOC_1 allows to obtain a closed-form solution for the optimal level of maintenance:

$$h_t^* = \frac{-\alpha R(y_t + R^{-1}y_{t+1}) - \alpha E_t(p_{t+1})H_{t-1}(1-\delta)^2\theta_1}{E_t(p_{t+1})(1-\delta)\theta_2 - R\kappa\theta_3}$$

Which is positive since both arguments are negative. The parameters θ_1,θ_2,θ_3 are such that: $\theta_1 = \frac{(\beta+\alpha)(1-\kappa R)}{1-\alpha}$, $\theta_2 = \alpha - \frac{(\alpha^2+\alpha\beta)}{1-\alpha}$, $\theta_3 = \alpha - \beta - \frac{(\alpha^2+\alpha\beta)}{1-\alpha}$. Also, the equation is increasing in the expectations of future house prices, given that the numerator increases (in absolute value) more than the denominator for increasing expectations. Eventually the optimal saving rate can be obtained by plugging h_t^* in the equation for a_t^*, while the optimal consumption can be obtained by further plugging h_t^* and a_t^* into the period t budget constraints.

Quality-specific pricing: Under this assumption, eq. (4.11) becomes:

$$c_t^* = \frac{\alpha}{1-\alpha}[R^{-1}(1-\delta)E_t(p_{t+1}(H_{t+1}(h_t))(H_{t+1}(h_t))) - \kappa](H_{t-1}(1-\delta) + h_t^*)$$

And its derivative with respect to the level of expectations is:

$$\frac{\partial c_t^*}{\partial E_t(p_{t+1})} = \frac{\alpha}{1-\alpha}\left[R^{-1}(1-\delta)\left(p'_{t+1}(H_{t+1}(h_t))H_{t+1} + p_{t+1}(H_{t+1}(h_t)\right)\right](H_{t-1}(1-\delta) + h_t^*) + ...$$

$$... + \frac{\alpha}{1-\alpha}[R^{-1}(1-\delta)E_t(p_{t+1}(H_{t+1}(h_t))(H_{t+1}(h_t)))]\frac{\partial h_t^*}{\partial E_t(p_{t+1})}$$

Bibliography

[1] Acharya, V., Bergant, K., Crosignani, M., Eisert, T., McCann, F. (2017) The Anatomy of the Transmission of Macroprudential Policy: Evidence from Ireland. Working Paper.

[2] Addison, J. T. (2016) Collective bargaining systems and macroeconomic and microeconomic flexibility: the quest for appropriate institutional forms in advanced economies *IZA Journal of Labour Policy* 5 (19).

[3] Adelino, M., Schoar, A., Severino, F. (2016) The Anatomy of the Transmission of Macroprudential Policy: Evidence from Ireland. *The Review of Financial Studies* vol. 29 no. 7

[4] Alessie, R., Lusardi, A., Van Rooij, M. (2011). Financial Literacy and Stock Market Participation. *Journal of Financial Economics*, vol. 101, n. 2, pp. 449-472.

[5] Alessie R., Brugiavini A., Weber G., (2005). Saving and Cohabitation: The Economic Consequences of Living with One's Parents in Italy and the Netherlands *NBER Working Paper*, N. 11079.

[6] Alessie, R., Hochguertel, S., van Soest, A. (2002). Household Portfolios in the Netherlands. In Guiso, L. and Jappelli, G. (Editors). Household Portfolios. Cambridge: The MIT Press.

[7] Altonji, J. G. and Devereux P. J. (1999) The Extent and Consequences of Downward Nominal Wage Rigidity. *NBER Working Paper* No. w7236.

[8] Attanasio, O., Weber, G. (1994) The UK Consumption Boom of the Late 1980's: Aggregate Implications of Microeconomic Evidence. *Economic Journal*, vol. 104, issue 427, pp. 1269-1302.

[9] Attanasio, O.P., Blow, L., Hamilton, R. and Leicester, A. (2009). Booms and Busts: Consumption, House Prices and Expectations. *Economica*, 76 (301), 20-50. .

[10] Barattieri, A., Basu, S. and Gottschalk, P. (2014) Some Evidence on the Importance of Sticky Wages. *American Economic Journal: Macroeconomics*, Vol. 6 (1), pp. 70-101.

[11] Bertola, G., Dabusinskas A., Hoeberichts M., Izquierdo M., Kwapil C., Montornès J. and Radowski, R. (2012) Price, wage and employment response to shocks: evidence from the WDN survey. *Labour Economics*, Vol. 19, pp. 783?791

[12] Bathia, K (1987). Real Estate Assets and Consumer Spending *The Quarterly Journal of Economics,* 102(2), 437.444.

[13] Best, M., Cloyne, J., Ilzetzki, E., Kleven, A. H. (2015) Interest rates, debt and intertemporal allocation: evidence from notched mortgage contracts in the United Kingdom. *Bank of England Staff Working Paper*, 543.

[14] Blanchard, O., Jaumotte, F. and Loungani, P. (2014) Labour Market Policies and IMF Advice in Advanced Economies during the Great Recession. *IZA Journal of Labor Policy*, 3(2).

[15] Boeri, T. and Jimeno, J. F. (2015) The elusive quest for labour market reforms. *VoxEU.org.*

bibitemlamport94 Breza, E., Chandrasekhar, A. G., McComick, T. H., Pan, M. (2020) Using Aggregated Relational Data to Feasibly Identify Network Structure without Network Data *American Economic Review*, forthcoming.

[16] Brown, M., Haughwout, A., Lee, D., Van der Klaauw, D. (2010) The Financial Crisis at the Kitchen Table: Trends in Household Debt and Credit. *Federal Reserve Bank of New York*, Staff Report.

[17] Brown, J., Ivkovic, Z., Smith, P., Weisbenner, S. (2008). Neighbors matter: Causal Community effects and Stock Market Participation. *The Journal of Finance*, vol. 63, n. 3, pp. 1509-1531.

[18] Browning, M., Gørtz, M. and Leth-Petersen, S. (2013). Housing Wealth and Consumption: A Micro Panel Study *The Economic Journal*, 123(568), 401-428. .

[19] Brueckner, J. K. (1994) The Demand for Mortgage Debt: Some Basic Results. *Journal of Housing Economics*, no. 3, pp. 251-262.

[20] Cai, J., De Janvry, A., Sadoulet, E. (2015). Social Networks and the Decision to Insure. *American Economic Journal: Applied Economics*, vol. 7, no. 2, pp. 81-108.

[21] Calvet, L. E., Campbell, J. Y., Sodini, P. (2007). Down or Out: Assessing the Welfare Costs of Household Investment Mistakes. *Journal of Political Economy*, 2007, vol. 115 (5), pp. 707-747.

[22] Calvet, L. E., Campbell, J. Y., Sodini, P. (2009). Measuring the Financial Sophistication of Households. *The American Economic Reviw*, vol. 99, no. 2, pp. 393-398.

[23] Calvo, G. A. (1983) Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, Vol. 12 (3), pp. 383-398.

[24] Cameron, A. C. and Miller, D. L. (2015) A Practitioner's Guide to Cluster-Robust Inference. *The Journal of Human Resources*, Vol. 50 (2), pp. 317-373.

[25] Campbell, J.Y. and Cocco, J.F. (2007). How Do House Prices Affect Consumption? Evidence From Micro Data. *Journal of Monetary Economics*, 54(3), 591-621.

[26] Card, D., Cardoso, A. R., Heining, J. and Kline, P. (2018) Firms and Labor Market Inequality: Evidence and Some Theory. *Journal of Labor Economics*, 2018, vol. 36 (1), pp. S13-S70.

[27] Carrol, C., Misuzu, O., Slacalek, J. (2011). How Large Are Housing and Financial Wealth Effects? A New Approach. *Journal of Money, Credit and Banking*, 43(1): 55?79.

[28] Cerrutti, E., Claessens, S., Laeven, L. (2017) The Use and Effectiveness of Macroprudential Policies: New Evidence. *Journal of Financial Stability*, vol. 28, pp. 203-224.

[29] Case, K.E., Quigley, J.M. and Shiller, R.J. (2005). Comparing Wealth Effects: The Stock Market versus the Housing Market. *B.E. Journal of Macroeconomics*, 5(1), 1-32.

[30] Chari, V. V., Jagannathan, R. (1988) Bank Panics, Information and Rational Expectations Equilibrium. *The Journal of Finance*, vol. 43, no. 3, pp. 28-30.

[31] Chetty, R., Friedman, J. N., Olsen, T., Pistaferri, L. (2011) Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records. *The Quarterly Journal of Economics*, vol. 126(2), pages 749-804

[32] Chiuri M.C., Jappelli T., (2003) Financial market imperfections and home ownership: A comparative study *European Economic Review*, 47(5), 857-875.

[33] Christiano, L. J., Eichenbaum, M. and Evans, C. L. (2005) Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy* Vol. 113, (1), pp. 1-45.

[34] Claessens, S. (2017). An Overview of Macroprudential Policy Tools. *Annual Review of Financial Economics*, Vol. 7, pp. 397-422.

[35] Cooper, D., Dynan, K. (2016) Financial market imperfections and home ownership: A comparative study *Journal of Economic Surveys*, (30)1, 34-55

[36] CPB (2018) Macro-economische Verkenningen 2019.

[37] Danzinger, L. and Neuman, S. (2005) Delays in Renewal of Labor Contracts: Theory and Evidence. *Journal of Labor Economics*, Vol. 23 (2), pp. 341-372.

[38] Deaton, A. (1992). Understanding Consumption. Oxford: Oxford University Press.

[39] De Fusco, A., Paciorek, A. (2017). The Interest Rate Elasticity of Mortgage Demand: Evidence from Bunching at the Conforming Loan Limit. *The American Economic Journal: Economic Policy*, vol. 9, no. 1, pp. 210-240.

[40] De Fusco, A., Johnson, S., Mondragon, J. (2019). Regulating Household Leverage. *The Review of Economic Studies*, Vol. 87 (2), pp. 914?958.

[41] De Nederlandsche Bank (2017). The Housing Market in Major Dutch Cities. *Occasional Study*, vol. 15-1.

[42] De Nederlandsche Bank (2019). *Overview of Financial Stability*, Autumn 2019

[43] De Nederlandsche Bank. *Annual Report.* 2009.

[44] Diamond, D. W., Dybvig, P. H. (1983) Bank Runs, Deposit Insurance and Liquidity. *The Journal of Political Economy*, vol. 91, n. 3, pp. 401-419.

[45] Diez-Catalàn, L. and Villanueva, E. (2014) Delays in Renewal of Labor Contracts: Theory and Evidence. *Banco de España Working paper*, no. 1431.

[46] Disney, R., J., Gathergood, J. and A. Henley (2010). House Price Shocks, Negative Equity and Household Consumption in the United Kingdom. *Journal of the European EconomicAssociation*, vol. 8(6), pp. 1179-1207.

[47] Duflo,E., Saez,E. (2012). Participation and Investment Decisions in a Retirement Plan: the Influence of Colleagues Choice. *Journal of Public Economics*, vol. 85, pp. 121-148.

[48] Dustmann, C., Ludsteck, J. and Schönberg, U. (2009). Revisiting the German Wage Structure. *The Quarterly Journal of Economics*, vol. 85, pp. 843?881.

[49] Elsby, M. W. L. (2009) Evaluating the economic significance of downward nominal wage rigidity. *Journal of Monetary Economics*, Vol. 124 (2), pp. 154-169.

[50] Elsby, M. W. L. (2019) How Prevalent is Downward Nominal Wage Rigidity? *Journal of Economic Perspectives*, Vol. 33, (3), pp. 185-201.

[51] Engelhardt, G.V. (1996). House Prices and Home Owner Saving Behavior *Regional Science and Urban Economics* , 26(3-4), 313-336..

[52] Fougère, D., Gautier, E. and Roux, S. (2018) Wage floor rigidity in industry-level agreements: Evidence from France. *Labour Economics*, Vol. 55, pp. 72-97.

[53] Financial Stability Forum (2001) Working Group on Deposit InsuranceProgress Report

[54] Fischer, S. (1977) Long-Term Contracts, Rational Expectations, and the Optimal Money Supply Rule. *Journal of Political Economy*, vol. 85, (1), pp. 191-205

[55] Friedman, M. (2014) Nobel Lecture: Inflation and Unemployment. *Journal of Political Economy*, 1977, Vol. 85, (3), pp. 451-72.

[56] Friedman, M., Schwartz, A. J. (1963) A monetary history of the United States and the United Kingdom. Chicago: University of Chicago press.

[57] Galati, G., Moessner, R. (2017). Macroprudential policy: a Literature Review. *Journal of Economic Survey*, vol. 27, no. 5, pp. 846-878.

[58] Gertler, M. and Trigari, A. (2009) Unemployment Fluctuations with Staggered Nash Wage Bargaining, *Journal of Political Economy*, Vol. 117 (1), pp. 38-86

[59] Gertler, M., Kiyotaki, N. (2015) Banking, Liquidity and Bank Runs in an infinite Horizon Economy. *The American Economic Review*, vol. 105, no. 7, pp. 2011-2043.

[60] Guiso, L., Haliassos, M., Jappelli, T. (2003). Households Stockholding in Europe: where do we stand and where do we go? *Economic Policy*, vol. 18, n. 36, pp. 123-170.

[61] Guiso, L., Jappelli, T. (2005). Awareness and Stock Market Participation. *Review of Finance*, vol. 9, n. 1, pp. 537-567.

[62] Haliassos, M., Bertaut, C. C. (1995). Why do so Few Hold Stocks? *The Economic Journal*, vol. 105, no. 432, pp. 1110-1129.

[63] Haliassos, M. Georgarakos, D., Pasini, G. (2014). Household Debt and Social Interactions. *Review of Finance*, vol. 15, no. 4, pp. 693-725.

[64] Hall, R. E. (2005) Employment Fluctuations with Equilibrium Wage Stickiness. *The American Economic Review*, Vol. 95 (1), pp. 50-65.

[65] Hansen, L. P., Singleton, K. J. (1983) Stochastic Consumption, Risk Aversion, and the Temporal behavior of Asset Returns. *Journal of Political Economy*, vol. 91, no. 2, pp. 249-265.

[66] Hijzen, A., Martins, P. and Parlevliet, J. (2019) Frontal assault versus incremental change: A comparison of collective bargaining in Portugal and the Netherlands. *IZA Journal of Labor Policy*, n. 11113.

[67] Hilber, C. A. L., Vermeulen, W. (2016). The Impact of Supply Constraints on House Prices in England. *Economic Journal*, vol. 126, pp. 358-405

[68] Hong, H., Kubik, J., Stein, J. (2004). Social Interactions and Stock Market Participation. *The Journal of Finance*, vol. 59, n. 1, pp.137-163.

[69] Iyer, R., Jensen, T., Johannesen, N., Sheridan, A. (2017) The Run for Safety: Financial Fragility and the Deposit Insurance. Working paper.

[70] Iyer, R., Puri, M. (2012). Understanding Bank Runs: the importance of Bank Depositors Relationship and Networks. *The American Economic Review*, vol. 102, n. 4, pp. 1414-1445.

[71] Juster, T., Lupton, J. P., Smith, J. P., Stafford, F. (2005). The Decline in Savings and the Wealth Effect. *The Review of Economics and Statistics*, 87(4), pp. 20?27

[72] Kaplan, G., Mitman, K., Violante, G.L. (2017) The Housing Boom and Bust: Model meets Evidence. *NBER Working Paper*, No. 23694.

[73] Kleven, H. J. (2016). Bunching. *Annual Review of Economics*, Vol. 8, pp. 435-464.

[74] Kleven, H. J., Waseem, M. (2013). Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan. *The Quarterly Journal of Economics*, Oxford University Press, vol. 128(2), pp. 669-723.

[75] Kim, Y. and Mastrogiacomo, M. (2020) The Effect of the Financial Assessment Framework on Mortgage Investments of Pension funds. Working paper.

[76] ECB (2009) Wage Dynamics in Europe: Final Report of the Wage Dynamics Network (WDN).

[77] Le Bihan, H., Montornès and J., Heckel, T. (2012) Sticky Wages: Evidence from Quarterly Microeconomic Data. *American Economic Journal: Macroeconomics*, Vol. 4 (3), pp. 1-32.

[78] Lee, M. J. (2005). Micro-econometrics for policy evaluation, program and treatment effects. Advanced Text in Econometrics, Oxford Scholarship Online.

[79] Mastrogiacomo, M., Van der Molen, R. (2015) Dutch Mortgages in the Loan Level Data. *DNB Occasional Studies*, vol 13 (4).

[80] Mayer, C., Pence, K., Sherlund, S. M. (2009). The Rise in Mortgage Deafults. *Journal of Economic Perspective*, vol. 23, no. 1, pp. 27-50.

[81] Mian, A., Sufi, A (2009). The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis. *The Quarterly Journal of Economics*, Volume 124, Issue 4, 1 November 2009, pp. 1449-1496.

[82] Mian, A., Sufi, A (2011). House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis. *The American Economic Review*, vol. 101, no. 5, pp. 2132-56.

[83] Mian, A., Sufi, A (2010). The Great Recession: Lessons from Microeconomic Data. *The American Economic Review: Papers and Proceedings*, vol. 100, pp. 51-56.

[84] Mian, A., Rao, K., Sufi, A. (2013). Household Balance Sheet, Consumption and Economic Slump. *The Quarterly Journal of Economics*, Volume 128, Issue 4, Pp. 1687?1726.

[85] Mian, A., Sufi, A., Trebbi, F. (2015). Foreclosures, House Prices, and the Real Economy. *The Journal of Finance*, 70, pp. 2587-2634.

[86] Mian, A., Straub, Sufi, A. (2020). Indebted Demand. Working Paper

[87] Mortensen, D. T. and Pissarides, P. A. (1994) Job Creation and Job Destruction in the Theory of Unemployment. *The Review of Economic Studies*, Vol. 61, (3), pp. 397-415.

[88] Muellbauer, J., Murphy, A. (1997). Boom and Busts in the UK Housing Market. *Economic Journal*, vol. 107, issue 445, pp. 1701-27

[89] Olivei G. and Tenreyro, S. (2010) Wage-setting patterns and monetary policy: International evidence. *Journal of Monetary Economics*, vol. 57, pp. 785?802

[90] Paiella, M. and Pistaferri, L. (2017) Decomposing the Wealth Effect on Consumption *The Review of Economics and Statistics*, 99 (4), 710-721.

[91] Piazzesi, M., Schneider, M. (2016). Housing and Macroeconomics. *NBER Working paper*, No. w22354.

[92] Ravi, D. and Zhu, N. (2006) Up Close and Personal: Investor Sophistication and the Disposition Effect. *Management Science*, vol. 52, pp. 726-740.

[93] Roberts, J. (1995) New Keynesian Economics and the Phillips Curve. *Journal of Money, Credit and Banking* Vol. 27 (4), pp. 975-84.

[94] Rouwendal, J. and Alessie, R. (2002). House Prices, Second Mortgages and Household Savings: An Empirical Investigation for the Netherlands, 1987-1994. *Timbergen Institute Discussion Paper*, , 074/3.

[95] Saez, E. (2012) Do Taxpayers Bunch at Kink Points? *American Economic Journal: Economic Policy*, Vol. 2, No. 3, pp. 180-212

[96] Schmitt-Grohé, S. and Uribe, M. (2016) Downward Nominal Wage Rigidity, Currency Pegs, and Involuntary Unemployment. *Journal of Political Economy*, Vol. 124, (5), pp. 1466-1514.

[97] Schularik, M., Taylor, A. M. (2012) Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008. *The American Economic Review*, vol. 102, no. 2, pp. 1029-61.

[98] Skinner J.S. (1988). Risky income, life cycle consumption, and precautionary savings *Journal of Monetary Economics*, vol. 22 (2), pp. 237-255.

[99] Skinner J.S. (1996). Is Housing Wealth a Sideshow? in Advances in the Economics of Aging. Chicago, University of Chicago Press, pp. 241-268.

[100] Smets, F. and Wouters, R. (2007) Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *The American Economic Review*, Vol. 97 (3), pp. 586-606.

[101] Stein, J. (1995) Prices and Trading Volume in the Housing Market: A Model with Down-Payment Effects. *The Quarterly Journal of Economics*, Vol. 110, No. 2, pp. 379-406.

[102] Suari Andreu, E. (forthcoming). House Price Shocks and Household Saving.

[103] Taylor, J. B. (1979) Staggered Wage Setting in a Macro Model. *The American Economic Review*, Papers and Proceedings of the 91st annual meeting of the AEA, Vol. 69 (2), pp. 108-113.

[104] Taylor, J. B. (2016) The Staying Power of Staggered Wage and Price Setting Models in Macroeconomics. *NBER Working Paper*, No. 22356.

[105] Thiel, J. (2020) Competition and Dynamic Pricing in Markets with Consumer Switching Costs. Working Paper.

[106] Van't Hof, A. (2017) Designing Macroprudential Regulation and Supervision Outside the Scope of the Banking Union: Lessons from the Netherlands and Ireland. *Journal of Banking Regulation*, vol. 18, pp. 201-212.

[107] Van Dijk (2019) Commercial and Residential Real Estate market liquidity. *PhD thesis* Universiteit van Amsterdam.

[108] Visser, J. (2016) What happened to collective bargaining during the great recession? *IZA Journal of Labor Policy* 5(9).

[109] Visser, J. (2019) CTWSS Database, version 6.1. Amsterdam: Amsterdam Institute for Advanced Labour Studies (AIAS), University of Amsterdam.

[110] Vissing-Jorgensen, A. (2003). Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structure. *NBER Working Paper* n. 8884.

[111] Yannis, B., Georgarakos, D. and Haliassos, M. (2008) Portfolio Inertia and Stock Market Fluctuations. *Journal of Money, Credit and Banking*, vol. 42 (4), pp. 715-742.