



THE UNIVERSITY *of* EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e.g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

This work is protected by copyright and other intellectual property rights, which are retained by the thesis author, unless otherwise stated.

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author.

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given.

DYNAMICS OF DEBT ACCUMULATION:
THREE ESSAYS IN APPLIED MACROECONOMICS

Alessia De Stefani

A THESIS PRESENTED FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY



THE UNIVERSITY OF EDINBURGH
SCHOOL OF ECONOMICS

2016

Declaration

I certify that this thesis was written by myself and is the result of my own work, unless clearly stated and referenced. This thesis has not been submitted for any other degree or professional qualification.

Alessia De Stefani

Abstract

Debt and credit markets played a crucial role in recent economic history. This thesis is composed of three chapters, each of which explores some drivers of private and public debt accumulation throughout the past decade. The first two chapters are directly linked, and study some behavioural determinants of household debt accumulation in the United States in the run-up to the 2007-2008 financial crisis. The third chapter takes a different perspective, and focuses on the political economy of fiscal reforms.

In the first chapter, I study whether the growth in US household debt ahead of the 2007-2008 financial crisis can be attributed to shifts in the distribution of personal income across the US population. The underlying theoretical mechanism is based on the idea that if individuals are concerned with status, rising income inequality within a given social group might lead its relatively poorer members to consume a larger proportion of their resources, due to a desire to emulate the consumption levels of richer individuals (Duesenberry [1949]; Frank, Levine and Dijk [2014]; Bertrand and Morse [2016]). I test this hypothesis by exploiting state-level variation in top incomes over time, following the methodology proposed by Bertrand and Morse [2016]. The results I present in this chapter challenge the status-emulation theory of consumer behaviour during the 2000s credit boom. I show that, between 1996 and 2007, only low and middle-income *homeowners* increased their expenditure and debt-to-income ratios as a response to an increase in income inequality in their state of residence. I also show that the growth in income inequality was strongly correlated with house prices growth, across US states and metropolitan areas. The positive correlation between inequality and household debt in the pre-crisis US might therefore be simply explained by the wealth and collateral effects experienced by low and middle-income homeowners living in areas where inequality was growing at the fastest rates.

The lifting of credit constraints due to rising house prices have been a major driver of household debt accumulation ahead of the 2007-2008 financial crisis (Mian and Sufi [2011]). However, this effect might have been coupled with a generalized optimistic belief that the growth in house prices was likely to continue in the future (Case, Shiller and Thompson [2012]).

The second chapter therefore tests whether consumers hold realistic expectations about the housing market, and whether this is a driver of their consumption and saving decisions. Using the Michigan Survey of Consumers, I show that American

households have heterogeneous expectations about the future of house prices, which systematically depend upon household characteristics, as well as upon the history of past house price realizations in the local area of residence. I also analyze individual-level forecast errors to show that house price expectations are biased and inefficient. Changes in individual forecast errors are predictable from past house price realizations in the local area of residence: in particular, forecast errors are positively correlated with recent price trends, and tend to become over-optimistic in good times, and over-pessimistic in bad ones. The predictability of forecast errors from public information available at the time the forecast was made suggests a violation of full-information rational expectations theory. This systematic bias in house price expectations matters because consumers make financial decisions on the basis of their house price beliefs. By exploiting an exogenous shift in housing sentiment, I show that when individuals expect the value of their properties to rise, they borrow against the anticipated increase in home equity.

The third and final chapter shifts the focus to the political drivers of public debt and deficits. Public debt crises often call for the intervention of international financial institutions, such as the International Monetary Fund (IMF). In this chapter, I introduce a new panel dataset on planned fiscal policy prescriptions included in all IMF loans between 2002 and 2012, and use it to study how domestic politics of recipient countries influence the content of IMF lending agreements. I show that IMF policy prescriptions depend strongly on domestic politics and that fiscal conditions are shaped by a political force often neglected in public choice literature: the threat of extra-parliamentary opposition, or civil unrest. Extra-parliamentary opposition (measured as a population's propensity to riot and demonstrate) significantly reduces the stringency of fiscal policy conditions attached to IMF loans. It also reduces the number of reforms in the realms of public employment and labor markets. These results suggest that fiscal policy has a strong political component even during circumstances when domestic politics are commonly assumed to cease to matter, as they do in IMF agreements. Also, they suggest that voting is not the only mechanism through which politics enters the technical realm of economic policy.

Lay Summary

Debt and credit markets have been driving the dynamics of the real economy in recent years. The 2007-2008 financial market crash led many developed economies into the Great Recession. The year 2009 marked the onset of a sovereign debt crisis at the periphery of the Eurozone, which threatened the very survival of the currency union. To a large extent, these dynamics were triggered by private or public debt accumulation, and from subsequent attempts at deleveraging. This thesis is composed of three self-contained chapters, each of which discusses some possible drivers of private and public debt accumulation throughout the past decade. The first two chapters are directly linked, and study some determinants of household debt accumulation in the United States in the run-up to the 2007-2008 financial crisis. The third chapter takes a different perspective, and focuses on public deficits and fiscal reforms.

In the first chapter, I test whether income inequality can be considered a main driver of household debt accumulation in the run-up to the 2007-2008 financial crisis. I challenge some results presented in previous literature, which suggest that the rise in US household debt ahead of the recent financial crisis could be attributed to the increase in income inequality occurring across the US population in the same years. According to this theory, inequality might have generated a desire for higher consumption and debt levels in low and middle-income households, due to the willingness to emulate the consumption levels of richer individuals. The results I present in this paper challenge this particular understanding of consumer behaviour during the 2000s credit boom. I show that the relationship between inequality and household debt in the run-up to the 2007 financial crisis could be simply explained by the role of house prices, which increased at higher rates in high-inequality areas.

The increase in house prices lifted credit constraints, since homeowners were allowed to borrow against the value of houses they *already owned* (Mian and Sufi [2011]). The role of rising house prices on household debt accumulation might however have been enhanced by a generalized optimistic belief that the growth in house prices was likely to continue in the future (Case, Shiller and Thompson [2012]).

In the second chapter I therefore examine how American households form house price expectations. I find evidence that expectations about the future of the housing market are heterogeneous across the population, and that they depend

strongly upon past house price realizations in the local area of residence. I also show that expectations are systematically incorrect: forecast errors are positively correlated with recent price trends, and tend to become over-optimistic in good times, and over-pessimistic in bad ones. The way people form house price expectations matters because these expectations directly affect households' financial decisions. When consumers expect house prices to increase in the near future, their mortgage leverage ratios expand accordingly. This implies that households might be borrowing against an *anticipated* increase in home equity.

The third chapter shifts the focus on public sector debt, particularly on the political drivers of fiscal reforms. I study how domestic politics influence a government's decision about whether or not to undertake fiscal consolidation, in the context of lending agreements with the International Monetary Fund (IMF). I present a dataset including all fiscal conditions agreed between the IMF and various country representatives between 2002 and 2012. I show that macroeconomic conditions are unable to explain all the variation in fiscal adjustment policies. In fact, international and domestic politics have non-negligible effects on the content of the lending agreements. In particular, I show that extra-parliamentary opposition (measured as a populations' propensity to riot, strike and demonstrate) significantly reduces the stringency of fiscal policy conditions attached to IMF loans. It also reduces the number of conditions in the realms of public employment, privatizations, and labor market reforms. These results suggest that economic policy, and in particular fiscal consolidation, has a strong political component even during circumstances when domestic politics are commonly assumed to cease to matter, as they do in IMF agreements. Moreover, they suggest that voting is not the only mechanism through which politics enters the technical realm of economic policy.

Acknowledgements

First and foremost I would like to thank my supervisors, Michèle Belot and Robert Zymek, for their advice and support throughout these four years. Working under Michèle's supervision has been a privilege. She has patiently guided me through several exploratory journeys into different fields of economics and always had an open door whenever I needed help, feedback or advice. I am also grateful for her trust in me during my endless tours around the globe, through which she has never lost confidence that I would eventually come back. Robert's continuous support and advice have proved invaluable. Always very generous with his time, he taught me how to think about a research project and provided me with deep theoretical insights, as well as with patient and thorough feedback on every detail of this dissertation. Robert has also efficiently guided me through some of my recent career choices, a role that has often seen him replying to my frantic e-mails over weekends and for which I am very grateful. I am also indebted to Mike Elsbey, for his invaluable feedback during the development of the first two chapters of this dissertation. Also, at different stages of my Ph.D., Ahmed Anwar, Aniko Biro, Steven Dieterle, Maia Guell, Ed Hopkins, Andy Snell and Ina Taneva have been an important source of advice and encouragement. I am grateful to my examiners, Liang Bai and Stephan Heblich, for their excellent feedback and suggestions.

My work has received the support of many others, who greatly contributed to the ideas presented in this thesis. I want to thank in particular Elisa van Waeyenberge and Ben Fine for inviting me to spend a semester at SOAS as a visiting research student. The intellectual environment at SOAS, as well as the interactions with its faculty and graduate students, deeply shaped my current research interests and enlarged my horizons as an economist. The first two chapters of this dissertation were possible thanks to the training I received during my internship at the Fiscal Affairs Department of the International Monetary Fund. I am particularly grateful to David Coady, Csaba Feher and Mauricio Soto for their advice, mentoring and friendship. Working at the Fund was a challenging and rewarding experience, as well as great fun. This research was possible thanks to the financial support I received from the Economic and Social Research Council and from the Edinburgh School of Economics.

I also want to thank my friends, in Edinburgh and abroad. I owe to my fellow Ph.D. students, in particular Aspasia Bisopoulou, Ibrahim Inal, Gerdis Mar-

quardt and Martina Vecchi for keeping up the spirits, always being available for a chat, and for bearing with me through the years. I am also grateful to the friends I had the fortune to meet over these four years in different cities and circumstances, in particular Sergi Cutillas, Gabrielle Flinn, Caterina Forti Grazzini, Johannes Langer, Luca Larcher, Dimitry Lisitsyn, Carmen Saldana, and Pablo San Martin. In different ways all of you have helped shaping parts of this research and contributed to my mental health over the past four years. A special thank you goes to Marta Arniani, Alberica Bazzoni, Ludovica Gazzè, Sarah Parolin, Camilla Pietrabissa and Francesca Tonelli for always being a home for me, albeit in multiple countries at the same time. Vanessa Ferdinand has been a continuous source of inspiration and adventures. Her support (and proofreading skills) have been crucial during several phases of this Ph.D. Last but not least, I would like to thank my family. Fausto De Stefani for his example in work ethics and for unconditionally supporting my choices even when they were keeping me away from home for too long. Dorina Giani, Stefania Giani and Valeria Milani for their support. Finally my mother, Daniela Giani, for showing me what is possible.

Table of contents

Introduction	13
1.Income Inequality and Household Debt: the Role of Home Ownership	17
1.1 Introduction	17
1.2 Data description	22
1.2.1 The Consumer Expenditure Survey (CEX)	22
1.2.2 The Panel Study of Income Dynamics (PSID)	23
1.2.3 The American Housing Survey	24
1.3 Consumption and debt	25
1.3.1 Methodology and descriptive statistics	25
1.3.2 Consumption: baseline	27
1.3.3 Consumption: the role of home ownership	29
1.3.4 Household debt	31
1.4 Income inequality and house prices	32
1.4.1 Methodology and descriptive statistics	32
1.4.2 Inequality and house prices: US states	33
1.4.3 Inequality and house prices: US metropolitan areas	35
1.5 Conclusions	36
Figures	38
Tables	41
Appendix A1.1: Regional clusters and standard errors	46
2.Waves of Optimism: House Price Expectations and Credit Cycles	47
2.1 Introduction	47
2.2 Empirical analysis of house price expectations	52
2.2.1 Data: expectations in the Michigan Survey	52
2.2.2 Results: determinants of individual expectations	53
2.2.3 Testing the rationality of expectations: methodology	56
2.2.4 Results: house price forecast errors	58
2.3 House price expectations and the credit cycle	61
2.3.1 The identification problem: IV strategy	61

2.3.2	Mortgage-level data	65
2.3.3	First stage: housing sentiment and elections	67
2.3.4	Second stage: expectations and mortgage leverage	71
2.3.5	Different types of mortgages	72
2.4	Conclusions	74
	Figures	76
	Tables	78
	Appendix A2.1: Control variables	88
3.Fiscal Adjustment and Social Unrest in IMF-Supported Programs		
		89
3.1	Introduction	89
3.2	Theoretical framework: conditionality as a two-level game	92
3.3	Data and methodology	96
3.3.1	Data on IMF conditionality	96
3.3.2	Costruction of conditionality variables	97
3.3.3	Methodology and explanatory variables	100
3.4	Results: fiscal conditionality and social unrest	103
3.4.1	Fiscal conditionality: macroeconomic determinants	103
3.4.2	Fiscal Conditionality and Social Unrest: Baseline	105
3.4.3	Robustness to selection	107
3.4.4	Mediating factors: social unrest and domestic politics	109
3.5	Social Unrest and structural conditions	111
3.5.1	Results	111
3.6	Conclusions	114
	Figures	116
	Tables	118
	Appendix A3.1: Cross sectional dependence	125
	Appendix A3.2: Interpolation of dependent variables	127
References		129

Introduction

Debt and credit markets played a crucial role in recent economic history. The macroeconomic implications of the 2007-2008 financial crisis spread well beyond the US, leading many developed economies into the Great Recession. More recently, the onset of a sovereign debt crisis in the periphery of the Eurozone in 2009 threatened the very survival of the European Union as a political entity. To some extent, these events are still unfolding as I write.

The work I present in this thesis is motivated by such occurrences. Much of the policy debate, during the 2007-2008 crisis and its aftermath revolved around macro-prudential policy, such as bank regulation mechanisms, optimal monetary or fiscal regimes, and the combination of these aspects with international financial flows. This thesis is an attempt to shed light on some micro-level drivers of the macroeconomic events unfolding in the past decade, drivers that have been somewhat overlooked in the general narrative of these crises. The first two chapters are directly connected, focusing on the drivers of household debt accumulation in the United States in recent years, while the third chapter takes a different perspective by analyzing the political economy of fiscal reforms. All three chapters focus on the dynamics of debt accumulation, because they attempt to shed light on the evolution of private and public debt levels over time.

This thesis also finds a coherent framework in its methodological approach. Traditionally, applied macroeconomic analysis has relied on time-series econometrics, country-level panels and aggregate data as a standard methodological toolkit. I follow instead a recent and quite innovative strand of literature, which exploits microdata, fine-grained geographic variation and techniques typical of the applied microeconomic literature to address macroeconomic questions. The reliance on micro-level variation was crucial in identifying the precise mechanisms through which housing, credit markets and household consumption were linked in turn of the 2007-2008 crisis (Mian and Sufi [2010]). The use of microdata not only permits to better isolate the channels through which macroeconomic outcomes may be linked to each other but, by focusing on individual behaviour, it also allows to disentangle these channels from aggregate macroeconomic shocks. Moreover, agent-level heterogeneity is rapidly becoming the norm in structural macroeconomic analysis. The applied analysis of the degree of heterogeneity in response to

aggregate shocks is the natural complement to this evolution in economic theory.

In the first chapter of this thesis I exploit household-level heterogeneity in survey data to study whether and to what extent changes in the distribution of personal income across the US population have been a direct driver of the increase in leverage ratio in the American economy ahead of the Great Recession. The underlying theoretical mechanism is based on the idea that if individuals are concerned with status, rising income inequality within a given social group might lead its relatively poorer members to consume a larger proportion of their resources, due to a desire to emulate the consumption levels of richer individuals (Duesenberry [1949]; Frank, Levine and Dijk [2014], Bertrand and Morse [2016], Carr and Jayadev [2014]). The within-state variation over time granted by the use of household surveys allows me to assess whether a rise in top incomes at the state/year level induces higher expenditure and debt levels in low and middle-income households living in the same state at the same point in time.

In contrast with earlier literature (Bertrand and Morse [2016]), I show that there is no evidence of a direct effect of top incomes on consumption or debt levels of the American middle class, in the time frame under study. However, I identify one social group that reacts more than other to increases in state-level inequality: low and middle-income *homeowners*. Given an increase in top incomes relative to median incomes, homeowners in the bottom 80th percentile of the income distribution increase their consumption and debt-to-income ratios significantly more than the comparable group of renters (for whom the effect of inequality is actually negative). This effect is particularly strong for mortgage debt, and is not explained by a change in the credit conditions at the national level, nor at the local level, since the inclusion of interest rates charged on mortgages does not affect the result. I show instead that rising income inequality is strongly correlated with house price growth, within regions. The positive correlation between inequality and household debt might therefore be more simply explained by the wealth effect and collateral effects experienced by low and middle-income homeowners living in areas where inequality was growing at the fastest rates.

The findings in this first chapter are relevant for the interpretation of the demand-side mechanisms behind the 2000s credit boom. They suggest that household debt accumulation prior to the financial crisis might not have been motivated by a real or perceived income loss, or by a status-comparison behavioural mechanism. Rather, the driver might have been the wealth and collateral effect associated with higher home valuations (Mian and Sufi [2011]; DeFusco [2015]).

The feedback mechanism between the lifting of credit constraints and the growth in house prices has been an important driver of household debt accumulation and household consumption, ahead of the Recession. However, this mechanism might also have been amplified by a generalized optimistic belief that the growth in house prices was likely to continue in the future (Case Shiller and Thompson [2012]; Adelino, Schoar and Severino [2016]).

The second chapter therefore tests directly whether consumers hold realistic expectations about the future of the housing market, and whether this is a driver of their consumption and saving decisions. I use the microdata contained in the Michigan Survey of Consumers, and exploit its variation along the lines of geography and time, to analyse how American households form house price expectations and how expectations depend upon the characteristics of their local area of residence. I show that consumers hold heterogeneous beliefs about the future of house prices, which largely depend on the history of past house price realizations in the local area of residence. Experiencing a state-level house price increase over the previous year leads households to forecast a price increase significantly higher, at the one-year horizon. Exploiting the panel component of the Michigan Survey, I also estimate individual-level forecast errors, and present evidence of a systematic extrapolative bias in consumers' house price expectations. Whenever house prices are increasing in their area of residence, individuals become over-optimistic about the future (and *vice versa*). The predictability of forecast errors from past house price realizations indicates a violation of full-information rational expectations theory. This systematic bias matters because consumers seem to be making financial decisions on the basis of their house price beliefs. By exploiting an instrumental variable technique based on an exogenous shift in housing sentiment, I show that when individuals expect the value of their properties to rise, the leverage ratios on their mortgages rises accordingly.

This result has relevant policy implications: in presence of a systematic bias in consumer house price expectations, house prices and mortgage markets may be subject to endogenous excess volatility (over-optimism in booms, and over-pessimism in busts). Then boom and bust cycles in the housing market may therefore be recurrent (Bordalo, Gennaioli and Shleifer [2016]) and pose substantial risks to macroeconomic and financial stability.

The third chapter shifts the focus to public debt and deficits, and their political drivers. The 2007-2008 financial crisis affected the public balance sheets of many developed economies. This was due in part to the direct cost of bank rescues,

and in part to the cost of the fiscal stimuli which many governments initially implemented in order to sustain aggregate demand in face of the credit crunch and the recession. These events, coupled with some pre-existing fiscal imbalances, triggered the onset of a public debt crisis in the periphery of the Eurozone, in 2009.

Public debt crises of this sort often call for the intervention of international financial institutions, such as the International Monetary Fund (IMF). The policy advice provided by these institutions is generally presented as being motivated by purely economic circumstances. In the third chapter, I introduce a new panel dataset on fiscal policy prescriptions included in all IMF loans between 2002 and 2012, and use it to study to what extent the domestic politics of recipient countries influence the content of these lending agreements. I show that macro-economic conditions are unable to explain all the variation in fiscal adjustment policies. In particular, fiscal conditions are shaped by a political force often neglected in public choice literature: the threat of extra-parliamentary opposition, or civil unrest. I show that a population's propensity to riot, demonstrate and strike significantly reduces the extent of public budget reforms a country is required to carry out as a part of the IMF agreements. All else held equal, two additional episodes of civil unrest in each of the three years prior to the agreement reduce fiscal consolidation targets and structural conditionality in the realms of welfare state, labor markets and public employment reform by a significant amount. The results are robust to controls for the economic circumstances of the country; to political variables which previous literature has shown to be relevant in determining IMF conditionality; and to some potential selection issues. This evidence suggests that fiscal consolidation has a strong political component even during circumstances when domestic politics are commonly assumed to cease to matter, as in the case of IMF agreements. Moreover, it suggests that voting is not the only mechanism through which politics enters the technical realm of economic policy.

This thesis is an attempt to shed light over some recent historical events which left behind many open questions for the field of macroeconomics. I see these chapters as the beginning of a personal research agenda, as well as modest attempts to reach a better understanding of the events unfolding in the past decade.

Income Inequality and Household Debt: the Role of Home Ownership

1.1 Introduction

A growing body of literature suggests that rising income inequality was at the root of the recent financial crisis (Van Treeck[2014]). This literature highlights the correlation between two prominent trends affecting household balance sheets over the past three decades: the growth of income inequality and the growth of household debt levels (Figure 1.1). The saving rates of American low and middle-income consumers, in particular, began declining as soon as inequality started to soar, in the middle 1980s (Rajan [2010]). This has led some authors to argue that poor and middle-income American consumers borrowed beyond their own capacity to repay in order to sustain consumption levels despite stagnating real incomes, generating fragility in the financial system (Rajan [2010]; Kumhof Ranciere Winant [2015]).²

This chapter attempts to ascertain whether rising income inequality can be considered a direct driver of the consumption choices of low and middle-income American consumers in the run-up to the 2007-2008 financial crisis. I focus in particular on testing the relative income hypothesis (Duesenberry [1949]), which received substantial attention in the public discourse as well as in academic circles (Frank, Levine and Dijk [2014]; Bertrand and Morse [2016]; Carr and Jayadev [2015]; De Giorgi Fredriksen and Pistaferri [2016]). This theory suggests that an increase in income at the top of the distribution could directly drive consumption choices at the bottom. If individuals are concerned with status, rising income inequality within a given social group will lead its relatively poorer members

1. This chapter was previously circulated as a working paper with title : “Debt, Inequality and House Prices: Explaining the Dynamics of Household Borrowing Prior to the Great Recession” Edinburgh School of Economics Discussion Paper Series, n.259 (2015).

2. The decade before the crisis is the time frame in US history when, for the first time since before the Great Depression, the top decile income share increased beyond 45% (Piketty Saez [2014]). During the same time frame, US households accumulated almost half of the debt outstanding at the onset of the crisis (Figure 1.2).

to consume a larger proportion of their resources, due to a desire to emulate the consumption levels of richer individuals. Low and middle-income consumers might therefore accumulate higher debt despite no real income change. Recent empirical literature finds support for this hypothesis, suggesting that inequality might have been a direct driver of household debt accumulation in the US prior to the 2007 financial crisis (Bertrand and Morse [2016]; Carr and Jayadev [2015]).³

The results I present in this chapter challenge this particular understanding of consumer behaviour during the 2000s credit boom. My results show that between the mid-1990s and 2007 there is no significant empirical relationship between the growth in income inequality and the growth in consumption or debt levels of low and middle-income American households. On the other hand, exploiting household-level heterogeneity in the data allows me to identify one particular social group that reacted strongly to changes in income inequality within this time frame: low and middle-income *homeowners*. Homeowners living in geographical areas where inequality increased at the fastest rates displayed a positive growth in consumption and debt levels, while renters living in the same state did not. My results also show that the growth in within-region income inequality was associated with a higher-than-average growth in *house prices*, in the same region. Therefore, the empirical relationship between inequality and household debt in the context of the pre-crisis US might be simply explained by the wealth and collateral effects experienced by homeowners living in high inequality areas.

My empirical approach relies on the use of state-level variation in top incomes over time to analyze how the consumption and debt levels of all households falling in the bottom eight deciles of the income distribution respond to changes in income levels of richer households living in the same state. In other words, I test whether when relatively richer households become richer, relatively poorer households residing in the same state tend to consume more out of their income. Using the Consumer Expenditure Survey (CEX), I show that given a 10% increase in top incomes relative to median incomes, homeowners in the bottom eight deciles of the income distribution increase their consumption of non-durable goods (unrelated to housing expenditure) by 4.6 percent more than the comparable group of renters. Since the net effect of top incomes on renters' consumption is negative, this estimation implies that a 10% increase in top incomes relative to median incomes is correlated with an absolute increase in low and middle-income home-

3. For a complete survey of the recent literature on this topic see Van Treeck [2014].

owners' non-housing consumption worth 0.7 percent. Using the Panel Study of Income Dynamics (PSID) I also show that inequality has a positive effect on homeowners' debt-to-income ratios. Given a 10% increase in top incomes, the debt-to-income ratios of homeowners rise by 6.4 percent more than those of the comparable group of renters, or 1.8 percent in absolute terms. This effect is not explained by a change in credit conditions, as the interest rates charged on mortgages were not significantly affected by changes in income at the top of the distribution. This evidence might instead be explained by the fact that the within-region growth in income inequality was positively correlated with the growth in house prices. Using the PSID and the American Housing Survey (AHS), I show that a 10% increase in top incomes relative to median incomes, between 1994 and 2007, was correlated with an average yearly increase in the self-reported value of homes worth about 0.5% per year across US states and close to 0.9% per year across the main metropolitan areas.

These findings are relevant for the interpretation of the demand-side mechanisms behind the 2000s credit boom. Earlier research has shown that, during this time frame, homeowners have exploited the collateral effects arising from housing wealth (Aron et.al [2010]; DeFusco [2015]) and that home-equity loans were used to finance consumption (Mian and Sufi [2009;2011]). While my results do not necessarily discount the psychological driver at the heart of the relative income hypothesis, they show that the empirical results in support of such a theory could be substantially understating the role of a reduction in borrowing constraints based on the increase in housing wealth. It should however be emphasized that I do not attempt to address causality concerns. The within-state change in top incomes is not exogenous to the consumption levels of poorer households living in the same state or to the change in house prices: all can be simultaneously determined by state-level time-varying unobserved factors (technological or fiscal policy shocks, for example). Despite these concerns, my results clearly indicate that owners responded to changes in state-level income inequality, while renters did not; and that changes in inequality were highly correlated with changes in house prices, within states and metro areas. For these results to be capturing unobserved state-level variation (and thus for the OLS estimate to be upwardly biased) it would be necessary that the unobserved shock were to change simultaneously top incomes and the consumption levels of poorer homeowners, leaving poorer renters unaffected. In other words, housing wealth seem to have played a pivotal role in explaining the link between inequality-debt cycle during the decade that preceded the 2007 financial crisis. This evidence provides an alternative explanation for

this empirical relationship, which insofar has been based on two major theories.

The first line of thought is that inequality has historically led towards an easing of the credit conditions applied to poor borrowers (Rajan [2010]). During the decade preceding the 2007-2008 financial crisis, the average interest rate on household debt fell considerably (Figure 1.4). This is also the decade during which income inequality in the US grew at the fastest rates since before the Great Depression (Piketty and Saez [2014]). Inequality might be therefore associated with particularly low interest rates. This theory is formally developed by Kumhof et.al. [2015], in a DSGE model which establishes a link between inequality, private debt, and financial crises. Assuming a decreasing marginal propensity to consume over the income ladder, top earners might be induced to invest (or lend) proportionally more of their resources, as their income increases. The influx of savings in the market lowers the real interest rate, inducing higher borrowing on the side of the most credit-constrained consumers and endogenously leading to widespread defaults. Kumhof et.al. [2015] present results that are able to replicate the long term dynamics of the US economy (both for the Great Depression and the Great Recession). However, the empirical evidence in Coibion et al. [2014] challenges this hypothesis; using microdata on bank credit originations, they find that inequality actually reduced credit provision to poor applicants across the US. They interpret their results as indicative that inequality functions as a screening device for lenders, who choose to lend to richer applicants whenever possible. I test whether the price of debt is driving my results, and find no evidence that changing top incomes affect mortgage interest rates in any significant way.

The second line of thought is related to the relative income hypothesis, originally developed by Duesenberry [1949] and most recently formalized by Frank, Levine and Dijk [2014]. This chapter is mainly a contribution to this second strand of literature. According to the theory of relative income, consumers compare their standard of living with those of their reference group of peers. If inequality increases within a reference group, individuals may consume more, given constant real income, in order to keep up with the consumption levels of richer individuals. The desire for positional goods might therefore explain why debt arises even in absence of a real income change. This hypothesis has been tested empirically mostly on aggregate data: Christen and Morgan [2005] find a strong effect of inequality on debt levels in the US pre-2000. Bowles and Park [2005] find that higher inequality is associated with longer working hours. More recently, Carr and Jayadev [2015] and Bertrand and Morse [2016] use microdata on the American

economy of the past two decades to test the relative income hypothesis, finding evidence in its support. All else held equal, when low and middle-income Americans have been exposed to higher levels inequality, they have been saving less and consuming more, especially on visible goods (Bertrand and Morse [2016]).

I borrow extensively from Bertrand and Morse's [2016] methodology. I however extend their analysis in several directions. First of all, I focus on disentangling the role of income inequality from the role of house prices, by studying specifically the decade preceding the 2007 financial crisis. Bertrand and Morse [2016] explicitly address whether their results might be driven by an increase in house prices, and they dismiss this potential confounding factor via two empirical tests. In particular, they show that splitting the sample between homeowners and renters does not indicate evidence of a statistically different effect of inequality between the two groups. Also, they show that homeowners were reacting more strongly to increases in inequality prior to the start of the housing boom (before 1995). The discrepancy between my results and theirs is explained by the fact that Bertrand and Morse [2016] do not explicitly address whether homeownership is driving their results *during* the housing boom. However, this is the time frame when the majority of US household debt was accumulated. Also, by using interaction terms and by focusing on different expenditure categories (housing VS non-housing consumption) I am able to show that in the decade preceding the crisis there was a sizeable difference in the response to changes in income inequality between homeowners and renters. Furthermore, I extend their analysis by analyzing the effects of income inequality on household debt levels and leverage ratios (using the PSID), which are not addressed in Bertrand and Morse's paper. Finally, in the second part of this chapter, I provide evidence of a strong empirical relationship between the growth rates of income inequality and of house prices across US states and metropolitan areas. Overall, the evidence I present in this chapter suggests that inequality was not a direct driver of household debt accumulation ahead of the 2007-2008 financial crisis (as suggested by Bertrand and Morse). Rather, the channel through which income inequality affects consumption is likely to be via the wealth and collateral effects originating from rising house prices, which are highly correlated with growth in income inequality within geographical areas.

In order to identify the trends between inequality, consumption, debt and house prices, I rely on three different household surveys. Section 1.2 describes each of them in detail. Section 1.3 presents the first part of the empirical analysis, or the relationship between inequality and consumption/debt accumulation. Section

1.4 provides evidence of the empirical link between inequality and house prices. Section 1.5 briefly summarizes the chapter.

1.2 Data description

The empirical analyses of this chapter require disaggregated data on consumption, wealth, and income of a representative sample of American households over a relatively long span of time. For this reason, I gather data from several population surveys: the Consumer Expenditure Survey, the Panel Study of Income Dynamics, and the American Housing Survey. These are well-known but complex data sources and their usage requires some discretionary choices. This section describes these choices in detail.

1.2.1 The Consumer Expenditure Survey (CEX)

The Consumer Expenditure Survey (CEX) is the most comprehensive American household survey on consumer behavior, and Bertrand and Morse [2016] use this dataset in their estimations. Therefore, this is a natural starting point for my analysis. The sample has a cross-sectional structure, is composed by about 6000 households per year and is nationally representative through sample weights. It consists of two main questionnaires for each interview in any given year (each household is required to respond to four interviews per year).

The first questionnaire is the Diary Survey: this collects data on weekly expenditures of frequently purchased items such as food at home, food away from home, alcoholic beverages, smoking supplies, personal care products and services, and nonprescription drugs. The second questionnaire, the Interview Survey, collects data on monthly expenditures for housing, apparel and services, transportation, health care, entertainment, personal care, reading, education, food, tobacco, cash contributions, personal insurance, and pensions. Both surveys therefore collect data on non-durables expenditure.

I use the Interview Survey as it provides a wider range of information on households' consumption behaviour, and it allows me to discriminate between housing and non-housing expenditure. In order to be consistent with Bertrand and Morse [2016], I exclude families who fail to respond to all four interviews and families with zero total consumption. The exclusion of families that do not respond to all

surveys takes a large toll on the sample size: I am left with about 1000 families per given year. However, this is needed in order to construct a yearly measure of expenditure, since the questionnaires are conducted on a quarterly basis, and respondents are only required to give information on expenditure in the previous *quarter*. Therefore, in order to provide a precise estimate of the yearly expenditure for any category of goods, it is necessary to sum the responses for the entire set of (four) interviews. Yearly figures are necessary because the empirical strategy relies on the allocation of households to certain deciles of the income distribution, based on their annual income. Also, this facilitates comparison with other survey variables used in this chapter, which are always expressed in annual terms. Moreover, the exclusion of households who fail to respond to all four interviews is a standard procedure when working with the CEX.⁴

Somewhat differently from Bertrand and Morse [2016], I rely on the aggregate consumption categories reported in the summary expenditure variables of the Interview questionnaire: total expenditure and housing expenditure. I do so for simplicity, and to allow for an easier replication of my results. My CEX sample is restricted to the interviews taking place from 1996 onwards, since these are the only waves available from the Bureau of Labour Statistics Website.

1.2.2 The Panel Study of Income Dynamics (PSID)

I use the PSID because it provides information on households' wealth and financial liabilities, unlike the CEX. While other American household surveys provide this information (the Survey of Consumer Finances, for example), the PSID is the only source that also grants access to information about the geographical area of residence of respondents, a crucial part of my empirical methodology. The PSID grew substantially over the years, and from an original sample of about 6000 households in 1968, it now stands at about 8500 American families being continuously interviewed. My sample is restricted to 1996-2011. I only take into account families reporting both a positive level of income and of consumption in a given year, and I am left with a sample of about 14000 household heads observed over time. The discrepancy between the panel of families in each cohort and of household heads is due to family spin-offs and drop outs (the PSID is structured to followed individuals, rather than families).

4. Also Blundell Pistaferri and Preston [2008] apply the same criteria and are left with about 8% of the original CEX sample.

1.2.3 The American Housing Survey

I will use the American Housing Survey as a source of information on house price changes experienced by individual households in US metropolitan areas. The main advantage of using the American Housing Survey (AHS) over other surveys describing the US housing stock, is its combination of a panel structure with the provision of a vast array of information on housing quality. Families are required to provide detailed information about their homes, including square-foot size, number of bedrooms/bathrooms, and recent renovations. This information allows one to take into account changes in housing quality which might have affected its value. The AHS surveys around sixty-thousand families per year, and alternates the year when it samples National data with years when it samples a subset of Metropolitan Statistical Areas (MSAs). However, the lowest geographical level identifiable in the National survey is the macro region (NE, NW, SE, SW), and this impedes a direct comparison with the CEX and PSID.

The Metropolitan survey, on the other hand, captures more fine-grained geographical information. It cycles through a set of 21 metropolitan areas, surveying each one about once every six years. Like the national survey, the metro survey is longitudinal. However, metro survey samples have been redrawn more often than the national samples, and this reduces the time spans where the longitudinal dimension applies. During 1996-2008, the metro surveys were conducted four times. This allows me to identify two sets of information on family-level home value change. The first set is composed by MSAs surveyed in 1996 and 2004: Atlanta, Cleveland, Hartford, Indianapolis, Memphis, Oklahoma City, St. Louis and Seattle.⁵ The second panel was collected in 1998 and 2007, and comprises Boston, Baltimore, Houston, Minneapolis, Tampa, and Washington DC. My dataset is therefore composed of 14 MSAs and two panels: the two samples interviewed in 1996-2004 and in 1998-2007.

The metro survey also samples about 60,000 individuals per year, but many were excluded from the analysis according to the following criteria. About 45% of the respondents are renters. I also exclude households that report negative or zero income. Moreover, not all households respond to both waves of the survey, and those who do not are naturally dropped from the analysis (as my main

5. The AHS sample also includes Sacramento (CA), but these observations are excluded from the empirical analysis since I have no information on some of the covariates (for example, the elasticity of housing supply) for this MSA.

dependent variable is the change in value of their primary residence). I also exclude households who changed residence between the two time periods and those who changed the size of their houses through additions (measured as the change in the number of rooms). This is because I am interested in measuring the change in home value experienced by a given household on the *same* property, throughout the years. Changes in residence would impede such comparison (since in each interview households only report the *current* home value), and substantial modifications of the structure of the house would have a similar effect (since the change in home value would not only capture the change in price *per se*, but also the change in *size*). Overall, I work on a sample of about 9000 households.

1.3 Consumption and debt

This section focuses on whether changes in the distribution of income between the mid 1990s and the 2007 financial crisis were related to American households' propensity to consume and borrow. I will use two data sources for this purpose: the Consumer Expenditure Survey and the Panel Study of Income Dynamics.

1.3.1 Methodology and descriptive statistics

Since the mid-1990s, the distribution of income became more unequal across the US: in real terms the bottom 60% of American households experienced a real income loss between 1999 and 2007 (Figure 1.3). Nevertheless, their propensity to consume increased. In particular, the expenditure-to-income ratio of the bottom four deciles of the income distribution grew by 15% over this time span, despite the real income loss (Figure 1.3). Debt-to-income ratios also increased disproportionately at the bottom of the income distribution. The change in debt-to-income ratios for the bottom quintile is close to 200% in eight years, while the second quintile experienced a 100% increase. Surprisingly, this effect is not due to higher access of lower-income households to home ownership: the home ownership rate for the bottom four deciles of the income distribution decreased by an average 10% over this time span. This suggests that the rising aggregate debt-to income-ratio reflects the intensive, rather than extensive, margin.

The trends described by Figure 1.3 imply that while inequality increased between the mid 1990s and the late 2000s, consumption inequality has not followed suit (as systematically documented by Krueger and Perri [2006], among others). All

quintiles increased expenditure and debt levels, and poorer households even more so.

In this section I study the relationship between the levels of low and middle-income households' debt and consumption and the levels of inequality in their state of residence. This choice is rooted in the theory of relative income, which predicts that people will try to keep up with the behaviour (consumption levels) of their reference group. A reference group is defined as the people a given family is more likely to know and/or interact with. Therefore, focusing on the geographical area of residence has the purpose of identifying reference groups at the finest level of detail which the publicly available data allows to reach.⁶

To study the empirical relationship between inequality and consumption, I will follow closely the empirical methodology suggested by Bertrand and Morse [2015] which is based on the following equation:

$$\log Y_{ist} = a + \beta_1 X_{ist}^I + \beta_2 \log(80thPercentileIncome)_{st} + \beta_3 Z_{st}^I + \chi_s + \psi_t + \varepsilon_{ist} \quad (1)$$

where $\log Y_{ist}$ is household expenditure or outstanding debt for family i in state s at year t .⁷ X_{ist}^I is a vector of family specific characteristics, namely: total family income; number of adults and children in household; age, race, gender and educational attainment of household head; and home ownership status.⁸ Z_{st}^I is a vector of controls for state-level time-varying characteristics which might be correlated both with inequality and consumption or debt. The inequality measure $\log(80thPercentileIncome)_{st}$ is defined at the state/year level; following Bertrand and Morse [2016] it is the average annual income of the top 20% of the state/year income distribution as defined by the Current Population Survey (CPS).⁹

6. The publicly available CEX and PSID datasets do not report any level of geographic detail other than the state of residence. One could, of course, think about other (and probably more relevant) reference groups: for example along the lines of gender, occupation, educational attainment, or race. Both the PSID as the CEX contain detailed information at the household level, and would allow for a subdivision of households along these lines. However the CEX and the PSID are not designed to be representative within states and state-level sample sizes can be rather small. The state-level measure of inequality is therefore the finest level of detail I can reach while respecting the structure of the data and its sampling design.

7. Income and expenditure variables are expressed in real terms. The CPI measure is local (state-level) as computed by Carrillo, Early, Olsen (2014).

8. While Bertrand&Morse[2016] control for family income by including categories for income thresholds every \$2000, I instead control for the actual measure of family income (annual gross income before taxes).

9. While I could compute income distribution based on the CEX or on the PSID, the CPS is more reliable when it comes to income distribution analysis due to its much larger sample size (about 60,000 surveys per year). The CPS on the other hand does not collect data on assets, liabilities or consumption.

Time trends, assumed to affect all states equally at the same point in time, are taken into account by year fixed-effects; and so are the time-invariant characteristics of each state due to the inclusion of state fixed-effect. The coefficient β_2 therefore measures the effect of within-state change in top incomes over time on the consumption levels of poorer households living in the same state.

Using the Current Population Survey to measure the deciles of the state/year income distribution also allows one to identify the minimum income threshold required for a CEX or PSID household to fall into the top two deciles of their state/year cell. If the household falls into this category, it is dropped from the sample. All remaining families (those below the 80th percentile of the state/year cell) are defined as “non-rich” households, and are the sample upon which I run the estimations. All models include sample weights, and the residuals are clustered at the state level, to account for the presence of a common random effect within states across families.¹⁰ The model in Equation (1) simply tries to answer the following question: if average top incomes rise in a given state/year, do poorer families in the same state/year spend more?

1.3.2 Consumption: baseline

Since the starting point of my analysis is Bertrand and Morse’s [2016] methodology, this section provides a comparison between my baseline results and theirs. The main challenge associated with this exercise is the discrepancy in time frame: while Bertrand and Morse’ dataset spans from 1980 to the late 2000s, mine only spans 1996 onwards, since these are the only waves of data publicly available from the Bureau of Labour Statistics.¹¹ Moreover, Bertrand and Morse do not provide regression estimates split over a subset of years. This prevents an exact comparison of my results with theirs. Nevertheless, they run a model very similar to Equation (1), where top incomes interact with dummy variables indicating each decade in their sample (1980s, 1990s, 2000s).¹² Therefore, Table 1.1 provides a set of estimations based on Equation (1) only for the years after 1999, in order to provide baseline coefficients that are as closely comparable as possible with the coefficients estimated by Bertrand and Morse in their model with interaction terms. The results are displayed in Column 1 of Table 1.1, showing that during

10. A discussion of the clusterization of standard errors in this context is provided in Appendix A1.2.

11. Also, the CEX is collected yearly; however, to facilitate comparison with the PSID, I only study the years for which also the PSID sample is available, namely 1996, 1997, 1999, 2001, 2003, 2005, 2007, 2009, 2011.

12. Bertrand and Morse[2016], in Column 1, Internet Appendix A3.

the 2000s inequality had no statistically significant effect on non-rich households' consumption. The coefficient is positive (+0.19), but not statistically different from zero. This might seem a surprising result: however this coefficient is almost identical to the coefficient found by Bertrand and Morse (associated with the interaction term between top incomes and the dummy for the 2000s).¹³ While in general they find strong evidence of the relationship between inequality and higher consumption in the bottom deciles of the distribution, they also find this relationship to be weaker during the 2000s. The interaction term between this decade and top incomes is positive (+0.21) but not significant.

Bertrand and Morse control for wealth effects on consumption by splitting the sample between homeowners and renters.¹⁴ In columns 2 and 3 of Table 1.1, I run the same test: in the subset of sample covering the years between 2000 and 2011, owners do not seem to react to increasing inequality more than renters. In fact, splitting the sample between homeowners and renters seems to indicate that renters respond more to changes in inequality, which is also consistent with some of the baseline results presented by Bertrand and Morse [2016].¹⁵ However, by allowing *all* coefficients (including time and state dummies) to differ across the two models, one cannot easily test whether the difference in the coefficient associated with top incomes between owners and renters is significant.

I run a more precise test of whether the two coefficients are statistically different from each other, by including an interaction between the inequality measure with a dummy variable indicating homeownership status (Column 4). This model constrains all coefficients to be the same across the two groups, except for the effect of homeownership. This test shows that the higher effect of inequality on renters' consumption levels, albeit positive (+0.209), is not statistically different from zero.

Table 1.1 has the purpose of confirming that I can broadly replicate some of Bertrand and Morse's [2016] results on the time frame under study. However, Table 1.1 also shows that a differential effect of inequality between owners and renters could not have been inferred from a simple comparison of the coefficients in Columns 2 and 3. Since the interactive model instead allows an evaluation of the statistical significance of such difference, I will use this model, rather than a split of the sample, in the following estimations.

13. Bertrand and Morse[2016], in Column 1, Internet Appendix A3.

14. Bertrand and Morse [2016], Table 5, Panel A, Columns 1 and 2

15. Ibidem.

1.3.3 Consumption: the role of home ownership

I study consumption as a function of income inequality using the Consumer Expenditure Survey. The point of this exercise is to study the drivers of household debt accumulation *ahead* of the 2007 financial crisis. This is the time frame during which the US experienced the highest growth rates of household debt in three decades.¹⁶ It is also the time frame when inequality grew at the fastest rate since the decade preceding the Great Depression (Piketty and Saez [2014]). Therefore the positive correlation between inequality and household debt growth holds during this time frame. This section attempts to identify whether such relationship is *direct* or mediated by other factors.

Therefore, Table 1.2 replicates the analysis presented in Table 1.1 focusing on the pre-crisis years: the estimation includes the late 1990s and excludes the years after 2008.

Column 1 shows that, for the overall population, the relationship between top incomes and non-rich households' consumption is positive but not significantly different from zero. This is similar to the result found in Column 1 of Table 1.1, albeit with a different magnitude. This is the first interesting result of this section, because it suggests that during the decade when the majority of US household debt was accumulated, there is no evidence of a direct relationship between regional growth in income inequality and changes in the consumption behaviour of low and middle-income households.

In Column 2, I test whether household-level heterogeneity can shed some additional light on this result. I estimate the same model of Column 1 with the addition of an interaction term between top incomes home-ownership status. The interaction between inequality and home-ownership status is positive and strongly significant (1% level). During this time frame, non-rich homeowners responded to rising inequality with a much higher increase in consumption than renters did. Specifically, the elasticity of response was +0.35; for every 10% increase in the income of top earners in a given state/year, non-rich homeowners' real consumption increased by 3.5 percent *more than renters*.

A concern with this first specification is that the differential effect of inequality on homeowners expenditure might mechanically reflect the increasing house prices faced by new homebuyers over this time period. If an increase in inequality

16. See Figure 1.2.

at the state level was correlated with the rise the value of residential housing (perhaps because rising inequality reflects skilled migration and/or gentrification dynamics) then the higher expenditure on the side of homeowners might simply reflect the additional cost imposed on new home buyers by these social and demographic changes occurring at the state level. Rental prices do not necessarily follow purchasing prices, and this might alone explain the differential effect between owners and renters. Since I have no way to isolate new homeowners from existing homeowners (because the CEX is a repeated cross section in which families are observed at most for one year and gives no indication of when a household bought its current home), I eliminate housing expenditure from total expenditure. The increase in consumption on the side of homeowners was not due to higher expenditure in housing. In fact, Column 3 shows that expenditure for shelter is not significantly affected by rising top incomes, and that in this respect owners and renters do not differ from each other. This suggests that if inequality was correlated with higher house prices, this affected both purchasing prices and rental prices alike. On the other hand, Column 4 shows that the elasticity of non-housing consumption to rising inequality is about 0.49 percent higher for non-rich homeowners than for renters. This effect does not capture aggregate income trends, as the effect is robust to the inclusion of controls for median incomes at the state level (Column 5). This latter result suggests that homeowners (and homeowners only) were increasing non-housing consumption when inequality was increasing in their state of residence.

However, this result might be biased by the fact that homeowners, on average, are richer than renters, so they might be closer (socially/geographically) to the richest people in their states of residence. Since I have no information on the precise geographic location of these families besides their state of residence, column 6 tries to address this concern by exploiting their position in the income distribution. I include a dummy for each different decile of the income distribution that non-rich households fall into (ranging from decile 1 to decile 7, as the top two deciles are excluded from the estimations). This barely affects the coefficient of the interaction term and does not change its significance. Overall, Column 6 shows that a 10% increase in inequality is associated with an increase in non-housing consumption about 4.6 percent larger for homeowners than for renters. Since the effect of inequality on renters' consumption is negative (-3.9 percent), this estimation implies that a 10% increase in top incomes relative to median incomes is correlated with an increase in non-rich homeowners' non-housing consumption worth 0.7 percent. This effect is small, but positive and significant.

1.3.4 Household debt

I study the relationship between top income and non-rich households' debt accumulation using the PSID because, unlike the CEX, it contains information about households' financial liabilities.

Column 1 of Table 1.3 shows that there is no significant relationship between non-rich households' debt levels and changes in top incomes during this time frame. This is consistent with the results presented in Table 1.2. However, non-rich *homeowners* strongly and significantly increased their debt-to-income ratios as a result of the increase in inequality in their state of residence. The elasticity of their response is about 0.64 percent higher than that of renters' (Column 2). This implies that for a 10% increase in top incomes relative to median incomes, non-rich homeowners increase their debt-to-income ratios by 1.8 percent, on average. This result is robust to controls for median incomes and for the average state-level mortgage interest rate charged to families below the 80th percentile of the income distribution.¹⁷ The differential effect of inequality on homeowners' debt-to-income ratios is mostly due to mortgage debt: Column 3 shows that inequality has no significant differential effect on consumer debt levels, between the two groups.

However, as suggested by Kumhof et.al. [2015], changes in inequality might generate an influx of savings on the market and therefore reduce the real interest rate. These results might therefore be driven by the price of mortgage debt: median interest rates on mortgages fell considerably in the US during the housing boom (Figure 1.4). However that individual-level interest rates on first mortgages are not significantly affected by changes in top incomes versus median incomes (Column 4).¹⁸

Interestingly, leverage ratios (measured as outstanding mortgages to house value) seem to be negatively affected by changes in top incomes (Column 5). A 10% increase in top incomes is correlated with a *reduction* in non-rich homeowners' leverage ratios worth 0.7 percent. This suggests that if mortgages were the main component of the increase in the debt-income ratio of non-rich home owners, they were more than offset by an increase in house prices.

17. The PSID includes information on the interest rate charged on first-mortgages at the family level. In these regressions I include a control for the state-level interest rate, computed as a weighted average of the interested rates reported by respondent families in the state/year cell.

18. Also, Figure 1.4 shows that the dispersion in mortgage interest rates across US states is relatively low, suggesting that the lending market is mostly influenced by the FED rates, rather than by regional characteristics.

This section shows that the relationship between income inequality and consumption/debt of low and middle-income households is likely to have been mediated by homeownership status, prior to 2008. In particular, only non-rich homeowners were responding to increases in inequality in their state of residence by increasing their consumption and debt levels. An implication of this evidence is that rising inequality might have been correlated with a wealth or collateral effect experienced by non-rich homeowners, in the form of rising house prices.

1.4 Income inequality and house prices

In this section, I test to what extent inequality has been related to the change in house prices across the US, in the decade preceding the 2007-2008 financial crisis. I analyse this effect by means of two different data sources: the PSID, at the level of states, and the American Housing Survey, at the level of metropolitan areas.

1.4.1 Methodology and descriptive statistics

Between 1996 and 2008, the average American family reported an increase in the value of their main residence between 5 and 10 percentage points every two years (Figure 1.5). The wealth increase suddenly stopped after 2008. However, this change in house prices was not homogenous across the US territory. Some of the most striking differences can be seen, for example, across metropolitan areas. As Figure 1.6 shows, the Washington DC metropolitan area experienced a *real* increase in house prices worth about 60% between 1998 and 2007, almost double the national average. The house price growth rate in Houston, Texas, during the same years, was instead only around 10%.

In order to estimate the relationship between inequality and the increase in housing wealth, I estimate the following model:

$$\Delta p_{igt} = a + \beta_1 X_{igt}^I + \beta_2 \Delta \log(80thPercentileIncome)_{gt} + \beta_3 Z_{gt}^I + \chi_g + \psi_t + \varepsilon_{igt}, \quad (2)$$

where Δp_{igt} is the change in the (self reported) value of housing assets of family i , in geographical area g , between year $t-n$ and year t . This model, unlike equation (1), is expressed in first-differences, as both the PSID and the AHS data sources have a panel structure (unlike the CEX) and the methodology does not rely on interaction terms. My dependent variable is the change in home valuation at the family-level between two time periods. Using a model in changes, rather than levels, also prevents me from capturing spurious correlation between variables. For consistency, all other regressors are also expressed in first-differences.

X_{igt}^I is a vector of family-specific characteristics which include the log of levels of family income; the change in income between $t - n$ and t ; age, race, educational attainment, marriage status and gender of the household head; number of children in the household. The inequality measure, $\Delta \log(80thPercentileIncome)_{gt}$, is the same as in other specifications (average income in the top two deciles of the state/year income distribution, measured by the CPS), but is expressed in changes (between $t-n$ and t). I also run some robustness checks using a more standard measure of inequality, using changes in Gini coefficients (also computed from the March CPS).

Z_{igt}^I is a vector of geographical area-specific characteristics which might affect house prices, namely: changes in median incomes; changes in the elasticity of housing supply, measured by Saiz [2010]; change in homeownership rates measured by the CPS; 10-year change in population size measured by the census; average change in interest rate on mortgages reported by PSID respondents between $t - n$ and t . All geographical-area specific measures (including inequality) are expressed at the state level in PSID estimations and at the metro area level in the AHS estimations. Finally, geographic and year fixed effects are included as usual. All regressions are weighted with the sample weights provided in the surveys, and errors are clustered at the level of geographical areas (states in Section 1.4.2 and metropolitan areas in Section 1.4.3).

I exclude households which changed ownership status between the two periods; those who changed residence; and those who did not move but changed the size of their house (measured as the number of rooms).¹⁹ This is because Δp_{igt} needs to reflect the change in wealth experienced by homeowners on the *same* house (which did not go through major improvements that might have substantially affected its value).

1.4.2 Inequality and house prices: US states

Table 1.4 studies the relationship between changes in inequality and changes in house values at the level of US states between 1999 and 2011 using the PSID.²⁰ The dependent variable is the change in the value of the main residence for

19. The measure for the number of rooms is only available in the American Housing Survey, and not in the PSID.

20. The panel starts in 1999 because I rely on the *changes* in house prices, rather than their levels.

homeowners who did not change residence (or homeownership status) between year $t-2$ and t .

At an aggregate level, a 10% increase in top incomes was correlated with an average state-level house price increase of about 0.3% over two years (significant at 1% level) between 1999 and 2011 (Column 1). When looking at this result in the pre-crisis period, the correlation is higher (0.7% in two years Column 2). These aggregate effects are rather small, but Columns 3-6 provide evidence that this effect is larger in the micro-level estimation that allows for household heterogeneity.

Column 3 shows that between 1999 and 2007, a 10% increase in top incomes relative to median incomes was associated with an increase in the value of a family's main residence worth about 1% over two years (significant at 5% level). This implies a yearly average change of about 0.5% (all else held equal). Column 4 shows that this relationship is robust to the use of a more conventional measure of inequality, the Gini coefficient: 10% change in the Gini coefficient correlates with a change in house prices worth 2.3% over two years. Both regressions take into account time-trends in income, by controlling for the change in median incomes at the state level, as well as state and year fixed effects.

The American public was already perceiving the burst of the housing bubble in 2007, as reported by Case Shiller and Thompson [2012]. Expectations on future house price growth were rapidly changing for the worse, and consequently house price growth was already slowing down before the crisis erupted in late 2007 (Figure 1.4). Columns 5 and 6 provide a robustness check, by focusing on the years between 1997 and 2005. The effect of inequality during this time span is even stronger: a 10% increase in top incomes is related to a change in house prices worth about 1.6 percentage points in two years, or an average of 0.77% per year (Column 5). This effect is robust to the inclusion of family level fixed-effects, to take into account household-specific time invariant characteristics (Column 6). Here, too, the coefficient is very close to the weighted OLS regression (0.15).

The elasticity of housing supply, as expected, has a generally negative effect on the change in house prices although this is not always significant. The change in mortgage interest rates (average at the state level) displays the expected negative coefficient, even if it is only significant at the aggregate level (Columns 1 and 2). This is likely to reflect the fact similar households tend to be subjects to similar credit conditions, across US states. Likewise, changes in the homeownership rates are never significant in the micro-level estimations.

1.4.3 Inequality and house prices: US metropolitan areas

Since the changes in top incomes are positively and significantly correlated with housing appreciation across US states, I use the American Housing Survey to test whether this correlations holds at the level of US metropolitan areas. The American Housing survey rotates its panels across metropolitan areas every 8 years on average. Therefore each family during this time period reports a change in house value at most once: between 1994 and 2006 for the first group of metropolitan areas, and between 1998 and 2007 for the second group. Column 1 shows the macro-level effect for all MSAs: on average, the elasticity between a 1% increase in top incomes and an increase in house prices is 0.96. The elasticity of housing supply has a negative coefficient, while positive changes in population display an elasticity of 0.15 on house price increase. Higher median incomes (in levels) are also positively correlated with house price increase (elasticity 0.3).

Columns 2-6 estimate the micro-level effects. The estimated effect of a 10% increase in top income relative to median incomes is 7 percentage points in eight years, or roughly 0.87% per year, on average (Column 2). Columns 3 and 4 split the sample between the two waves of MSAs; the first (households interviewed in 1996 and in 2004) and the second (households interviewed in 1998 and in 2007). The estimated effect of inequality on house prices differs substantially between the two columns. The first wave of MSAs (Column 3) displays an elasticity of 2.4, implying that a change in top incomes relative to median incomes had a more than proportional effect on house prices during this time frame: a 10% increase in top incomes was associated with an overall increase in house prices worth 24 percentage points in eight years, or an average *yearly* increase in prices worth about 3%.

However, the estimated coefficient on top income for the second wave of MSAs (Column 4) is only 0.8% per 10% increase in top incomes (roughly 0.1% per year, on average). The positive effect of changes in inequality on house price increase is confirmed when using Gini coefficients at the MSA level (Columns 5 and 6). Again, the difference between the two waves is substantial: a 10% increase in Gini coefficients was associated with a house price increase of about 8% for the first group of MSAs, and only 1.2% for the second group. The differential in coefficients between these two waves of MSAs may cause some concern. However, the samples are composed of different metropolitan areas. In particular, the second group has an outlier in Houston (TX), which experienced house price increases well below the US average in this time frame, and is widely

regarded to have been a peculiar case among US metro areas during the boom.²¹ Moreover, this second wave of interviews was conducted when the price slowdown had already started (at the end of 2007).

Overall these results indicate that in the pre-crisis US, a regional growth in income inequality was strongly correlated with an higher than average growth in house prices.

1.5 Conclusions

The relative income hypothesis states that income inequality should be negatively correlated with the saving rates of low and middle-income households (Frank Levine and Dijk [2014]). This is because income inequality might shift the preferences of “non-rich” consumers, who will attempt to emulate the consumption levels of richer individuals and accumulate more debts in the process. Recent literature provides empirical evidence in support of this hypothesis, suggesting that income inequality might have been a direct driver of US household debt accumulation ahead of the 2007 financial crisis (Bertrand and Morse [2016]; Carr and Jayadev [2015]).

I test this theory on the years when the majority of US household debt was accumulated, between the mid 1990s and 2007. I find that the relationship between inequality and consumption holds only for a particular category of consumers: poor and middle-income *homeowners*. This group exhibited a positive expenditure reaction to increases in income inequality within their states of residence, especially with respect to non-housing expenditure. Non-rich homeowners also accumulated more mortgage debt, in response to increasing top incomes. However, their leverage ratios (mortgage to house value) did not increase, nor did the interest rates they paid on their debt.

This evidence points to a relationship between income concentration and house price growth across the US. I find empirical support in favour of this hypothesis. Exploiting both geographical and time variation across US states and metropolitan areas, I present evidence of a positive within-region correlation between inequality and the increase in house prices in the decade preceding the financial crisis of 2007-2008.

21. Houston, due to its large supply of land and permissive regulations, reacted to demand growth through higher construction, not price increase, largely avoiding the boom and bust dynamic which other cities experienced (FED, 2008). Excluding Houston from the regression for this set of MSAs yields a coefficients on top incomes of 0.95, significant at 1% level.

These results shed some light on the demand-side mechanism driving the 2000s credit boom. The relative income mechanism seems empirically indistinguishable from a more canonical wealth effect, against which households might have decided to borrow (Mian and Sufi [2009;2011]). An alternative and complementary interpretation is that households were subject to an increase in collateral availability, as suggested by Aron et.al [2010] and DeFusco [2015].

This analysis could be extended by studying *why* income inequality is correlated with rising house prices. The link between inequality and house prices has so far received little attention in the literature.²² While a large part of the housing boom taking place ahead of 2007 should probably be attributed to credit market liberalization (Mian and Sufi [2009]; Favara and Imbs [2015]; Jordà et.al [2015a]), the correlations I present here suggest that demand-side factors may have played an independent role in the dynamics of the real estate market during the boom years.

Overall my results suggest that the effect of inequality on consumer credit in the decade preceding the 2007 crisis was mediated by the role of house prices. In other words, what has largely been considered imprudent behaviour on the side of the weakest American consumers (borrowing beyond their own capacity to repay) might as well have been the result of a widespread illusion: the belief that the value of real estate would keep on growing (or at least hold its value indefinitely).²³

This changes the narrative, shifting the blame for the post-2008 recession from “poor” consumers to poor regulators. It suggests that if house prices were not to reach unsustainable levels, as a result of better urban planning and more regulated credit markets, it might have been possible to mitigate the effects of the Recession which followed the housing boom.

22. Matlack & Vigdor (2008) provide an analysis of rental prices in the US between 1970 and 2000, using census microdata. They find that in markets with low vacancy rates, increases in top incomes relative to the median imply a significant increase in the rental price per room paid by families at the bottom of the distribution. Maattanen and Tervio(2014) in a partial equilibrium model find negative effects of inequality on house prices, except in a segment of housing very close to the top of the distribution. However, their model is based on matching approach: each household holds one house, and there is no migration/foreign investment/speculation on the housing market. Zhang et.al [2016] find that income inequality correlates with an increase in the price-to-income ratios across Chinese cities.

23. On the role of beliefs in the run-up to the 2007-2008 financial crisis see, for example, Piazzesi and Schneider [2009] Case Shiller and Thompson [2012] or Adelino Schoar and Severino [2016].

Figures

Figure 1.1: Correlation between country/year growth rates of income inequality and country/year growth rates of household debt-to-GDP ratios. OECD countries, 1978-2011. Sources: Piketty and Saez[2014]; BIS Statistics.



Figure 1.2 Household Debt-to-GDP ratio, United States 1978-2011. Source: BIS Statistics.




Figure 1.3: Percentage change in selected indicators by quintile of the income distribution. United States 1999 to 2007. Source: PSID.

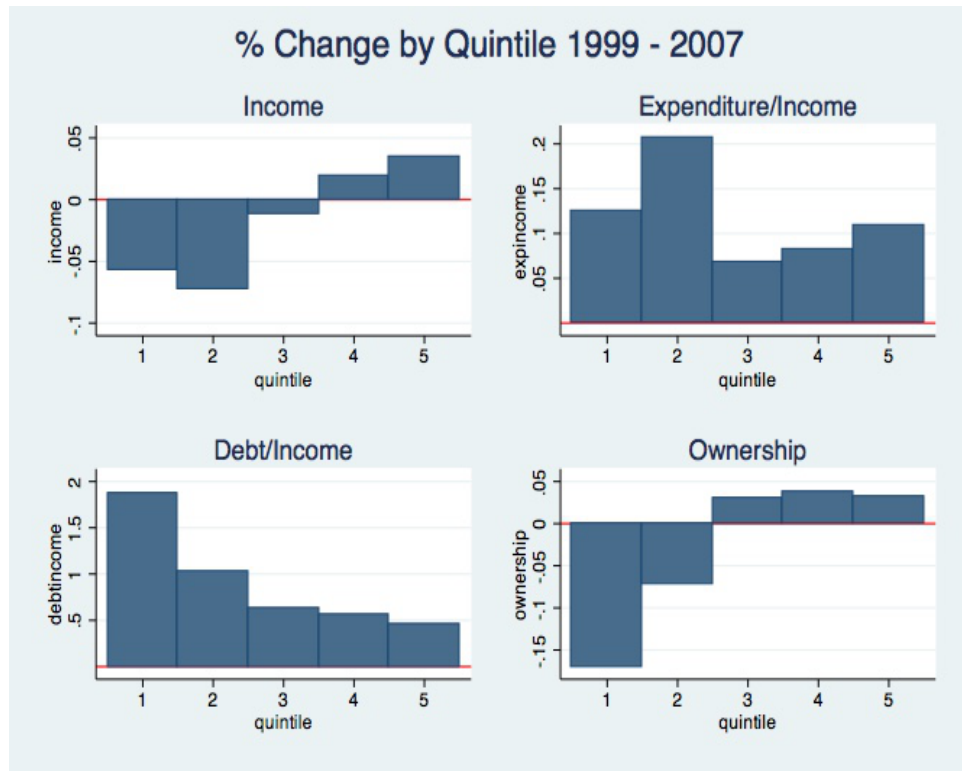


Figure 1.4. Evolution of average interest rates on first mortgages reported by PSID respondents over time (United States, 1996-2011). Source: PSID.




Figure 1.5. Average reported change in home value VS previous interview by year. United States 1997-2011. Source: PSID.

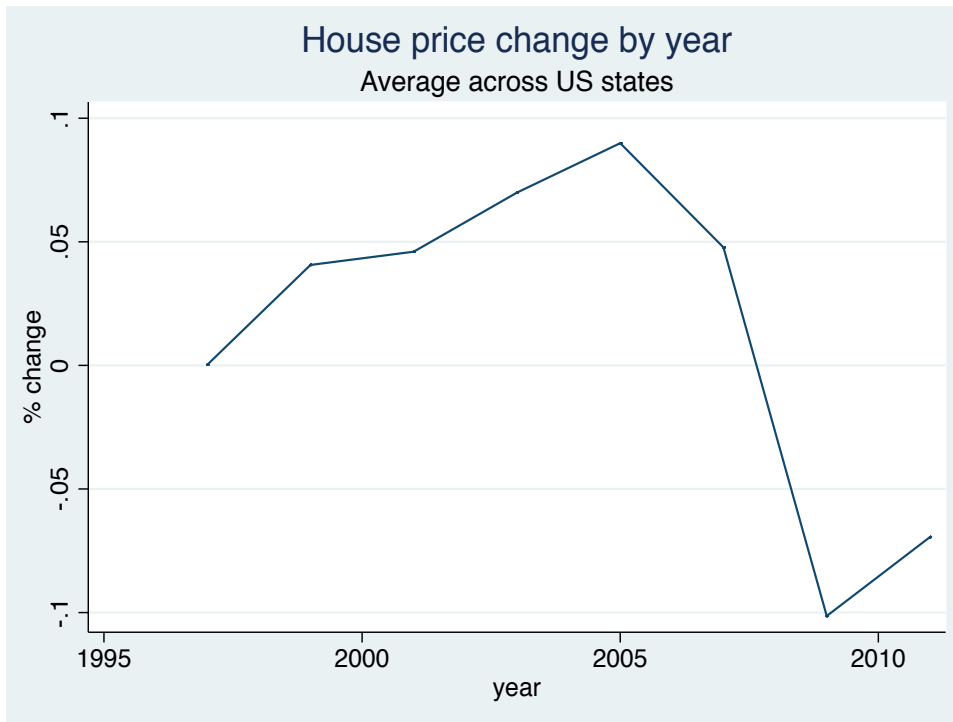



Figure 1.6. Average reported change in home value VS previous interview. Selected US metropolitan areas, 1996-2004 and 1998-2007. Source: American Housing Survey.



Tables

Table 1.1. Top income levels and bottom 80th percentile's consumption:2001-2011

VARIABLES	(1) Consumption	(2) Consumption Owner	(3) Consumption Renter	(4) Consumption
Top 20% Income	0.194 (0.170)	-0.035 (0.178)	0.674* (0.339)	0.209 (0.224)
Owner	0.195*** (0.026)			0.449 (2.473)
Top20*Owner				-0.022 (0.211)
State FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Family Controls	Yes	Yes	Yes	Yes
Constant	3.671* (1.998)	6.624*** (2.057)	-2.389 (4.035)	3.502 (2.652)
Observations	4,471	3,056	1,415	4,471
R-squared	0.337	0.308	0.351	0.337

Source: Consumer Expenditure Survey, 2001 to 2011. OLS regression. The sample is restricted to households below the 80th percentile of the state/year cell. The dependent variable is the logarithm of yearly total expenditure at the family level. All variables are in real terms, with CPI scaled at the State level (1996=1). Top 20% Income is the average income of families falling in the top 20% of the income distribution in a given state/year, computed from the March CPS. Family controls include a logarithm of income; age of head and its squared; sex, marital status, race and educational attainment of head; number of children in HH. Sample weights from the CEX are included. Errors are clustered at the State level.

*** p<0.01, ** p<0.05, * p<0.1

Table 1.2. Top income levels and bottom 80th percentile's consumption: 1997 to 2007.

VARIABLES	(1) Consumption	(2) Consumption	(3) Consumption Housing	(4) Consumption Non Housing	(5) Consumption Non Housing	(6) Consumption Non Housing
Top 20 Income	0.074 (0.170)	-0.185 (0.181)	-0.139 (0.232)	-0.203 (0.198)	-0.290 (0.201)	-0.387* (0.222)
Top20*owner		0.359*** (0.119)	0.099 (0.125)	0.492*** (0.121)	0.497*** (0.120)	0.459*** (0.133)
Owner	0.225*** (0.021)	0.230*** (0.018)	0.084** (0.023)	0.313*** (0.022)	0.312*** (0.022)	0.274*** (0.023)
Median Income					0.430 (0.312)	0.485 (0.325)
State FE	yes	yes	yes	yes	yes	yes
Year FE	yes	yes	yes	yes	yes	yes
Family Controls	yes	yes	yes	yes	yes	yes
Decile FE	no	no	no	no	no	yes
Observations	5,528	5,528	5,528	5,528	5,528	5,528
R-squared	0.329	0.330	0.239	0.353	0.353	0.369

Source: Consumer Expenditure Survey, 1996 to 2007. OLS regression. The sample is restricted to households below the 80th percentile of the state/year cell. The dependent variable is the logarithm of yearly total expenditure at the family level in columns 1-2; the log of housing expenditure in column 3; the log of non-housing expenditure (calculated as a residual) in cols 4-6. All variables are expressed in real terms (state-level CPI, 1996=1). The variable "owner" is centered in columns 2-6. Top 20% Income is the average income of families falling in the top 20% of the income distribution in a given state/year, computed from the March CPS. Family controls include a logarithm of income; age of head and its squared; sex, marital status, race and educational attainment of head; number of children in HH. Sample weights from the CEX are included. Errors are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1

Table 1.3. Top income levels and bottom 80th percentile's financial liabilities.

VARIABLES	(1) Debt/Income	(2) Debt/Income	(3) Consumer Debt/Income	(4) Interest rate	(5) Leverage
Top 20 Income	-0.098 (0.075)	-0.464*** (0.113)	-0.766 (0.698)	0.037 (0.057)	-0.071* (0.041)
Top20*Owner		0.646*** (0.215)	0.771 (0.852)		
Owner	0.596*** (0.039)	-7.007*** (2.513)	-10.277 (10.008)		
Median Income	0.045 (0.091)	0.030 (0.092)	-0.620 (0.578)	0.030 (0.064)	-0.123*** (0.040)
Average Interest	0.001 (0.049)	0.003 (0.049)	-0.019 (0.436)		
Constant	1.016 (0.993)	5.473*** (1.659)	12.414 (9.509)	1.544** (0.752)	2.596*** (0.582)
State FE	yes	yes	yes	yes	yes
Year FE	yes	yes	yes	yes	yes
Family_Controls	yes	yes	yes	yes	yes
Observations	41,742	41,742	41,743	15,678	24,878
R-squared	0.276	0.277	0.113	0.333	0.383

Source: Panel study of Income Dynamics, 1996-2007. OLS regression. The sample is restricted to households below the 80th percentile of the state/year cell. The dependent variable in columns 1-2 is outstanding debt to income ratio; in column 3 is non-mortgage debt; in column 4 is interest rate charged on the main mortgage; in column 5 is leverage (mortgage outstanding/house value). All variables are expressed in logs and are in real terms, with CPI scaled at the State level. Top 20% Income is the average income of families falling in the top 20% of the income distribution in a given state/year, computed from the March CPS. Family controls include age of head and its squared; sex, marital status, race and educational attainment of head; number of children in HH. Sample weights are included in all columns. Errors are clustered at the state level.

*** p<0.01, ** p<0.05, * p<0.1

Table 1.4 Top income levels and house prices: US states.

VARIABLES	(1) Average Delta Price 1999- 2011	(2) Average Delta Price 1999-2007	(3) Delta House Price 1999-2007	(4) Delta House Price 1999-2007	(5) Delta House Price 1999-2005	(6) Delta House Price 1999-2005
Δ Top Income	0.030*** (0.004)	0.069*** (0.004)	0.107** (0.052)		0.161*** (0.051)	0.152*** (0.045)
Δ Median Inc	0.135*** (0.005)	0.031*** (0.005)	0.121 (0.079)	0.211** (0.087)	0.141 (0.099)	0.088 (0.100)
Elasticity	-0.004 (0.008)	0.002 (0.009)	-0.019*** (0.002)	-0.019*** (0.002)	0.006** (0.003)	0.105 (0.173)
Δ Interest rate	-0.093*** (0.005)	-0.078*** (0.005)	-0.069 (0.070)	-0.070 (0.072)	-0.043 (0.079)	-0.040 (0.086)
Δ Ownership	0.022*** (0.008)	-0.039*** (0.008)	-0.103 (0.099)	-0.098 (0.097)	-0.099 (0.101)	-0.067 (0.132)
Δ Gini				0.235** (0.092)		
State and Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Family Controls	No	No	Yes	Yes	Yes	Yes
Family FE	No	No	No	No	No	Yes
Observations	52,067	37,428	13,273	13,273	10,876	11,038
R-squared	0.565	0.427	0.045	0.045	0.053	0.022
Number of familyID						4,223

Source: PSID, 1999 to 2011. OLS regression. The sample is restricted to households below the 80th percentile of the State/year cell. House Price levels are in real terms, with CPI expressed at the State area level. Delta house price is the year-on-year change in the value of house. Top incomes and median incomes and Gini coefficients (in changes) are calculated from the March CPS. Elasticity is the measure of elasticity of housing supply available from Saiz (2010). The change in mortgage interest rates is calculated at the state/year level from the PSID. Ownership rates for families falling in the bottom 80th percentile of the income distributions are also calculated from the PSID. Family level controls include age, education, race, sex, marriage status of the household head; also log of income and number of children at the family level. Sample weights from the PSID included in columns 3-5, col 6 includes family fixed-effects. The dependent variable in col. 1-2 is the average change in house prices at the State level; in columns 3-5 is the family-level change in house prices. Errors are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1

Table 1.5 Top income levels and house prices: US metro areas

VARIABLES	(1) Average Delta Price	(2) Delta House Price	(3) Delta House Price 1996-2004	(4) Delta House Price 1998-2007	(5) Delta House Price 1996-2004	(6) Delta House Price 1998-2007
Δ Top Income	0.961*** (0.012)	0.703* (0.388)	2.432*** (0.075)	0.084* (0.038)		
Δ Median Income	0.647*** (0.013)	0.751*** (0.208)	-1.696*** (0.131)	-2.316*** (0.057)	1.078*** (0.136)	-2.253*** (0.034)
Elasticity	-0.036*** (0.002)	-0.098 (0.062)	0.253*** (0.011)	-0.905*** (0.009)	-0.076*** (0.008)	-0.905*** (0.009)
Δ Population	0.150*** (0.003)	0.219** (0.076)	0.059*** (0.008)	0.546*** (0.008)	0.100*** (0.008)	0.551*** (0.009)
Log Median Inc.	0.300*** (0.019)	-0.118 (0.382)				
Δ Gini					0.808*** (0.025)	0.126* (0.057)
Area and Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Family controls	No	Yes	Yes	Yes	Yes	Yes
Observations	17,601	6,657	4,929	1,728	4,929	1,728
R-squared	0.925	0.133	0.019	0.109	0.019	0.109

Source: American Housing Survey, 1996 to 2007. WLS regression. The sample is restricted to households below the 80th percentile of the MSA/year cell. House Price levels are in real terms, with CPI expressed at the Metro area level. Delta house price is the change in the value per room reported by a panel of families interviewed between 1996-2004 and 1998-2007 in 14 metro areas. Top incomes and median incomes and Gini coefficients (in changes) are calculated from the March CPS. Elasticity is the measure of elasticity of housing supply available from Saiz(2010). Family level controls include age, education, race, sex, marriage status of the household head; also log of income and number of children at the family level. The variation in population is calculated as the 10 year average from Census. Area Fixed effects include a dummy for macro geographical areas: northeast (Baltimore, Boston, Hartford, Washington DC); South (Atlanta, Memphis, Oklahoma City, Tampa, Houston TX); Midwest (Cleveland, Indianapolis, Minneapolis, St.Louis). The dependent variable in column 1 is the average change in house prices at the MSA level; in columns 2-6 is the family-level change in house prices. Column 1 and 2 take into account both waves (change between 1996-2004 and 1998-2007) and all SMSAs. Columns 3 and 5 only restrict the analysis to the change occurring between 1996 and 2004 for the first wave of MSAs: Atlanta, Cleveland, Hartford, Indianapolis, Memphis, Oklahoma City, Seattle, St.Louis. Columns 4 and 6 restrict the analysis to the second wave, between 1998 and 2007: Baltimore, Boston, Houston, Minneapolis, Tampa and Washington DC. Each observation is weighted using the sample weights included in the AHS survey. Errors are clustered at the MSA level. *** p<0.01, ** p<0.05, * p<0.1

Appendix A1.1: Regional clusters and standard errors

In Equation (1) I define a model of the type:²⁴

$$(A1) Y_{ist} = \gamma_s + \lambda_t + \beta X_{st} + \epsilon_{ist}$$

Since the error term in (A1) reflects the idiosyncratic variation in outcomes across people, states, and time, some of this variation is likely to be common to all individuals residing in the same state at the same time (due to unobserved state/time shocks). The model (A1) can be written more precisely as:

$$(A2) Y_{ist} = \gamma_s + \lambda_t + \beta X_{st} + v_{st} + \eta_{ist}$$

Where v_{st} are state-year time shocks. With only two states and two time periods, there would be no way to distinguish the effect of X_{st} from these state/year shocks. The models presented in this chapter rely on 51 states or 16 metropolitan areas and therefore it is reasonable to assume that v_{st} will average out to zero as the number of groups (regions/years) increases (Angrist and Pischke [2008]).

However in (A2) the problem is not just cross-sectional, but has also a time dimension: v_{st} and \bar{v}_s (the average of v_{st}) are likely to be serially correlated (as regional shocks extend their influence over time).

The simplest method adopted to take this problem into account is to cluster errors at the state level, rather than at the state/year level, as would be appropriate if the problem related only to cross-sectional dependence (Bertrand et.al [2004]). Clustering at the state, rather than the state/year level, allows for residual correlation within clusters, including the time series correlation $v_{st} - \bar{v}_s$. Angrist and Pischke [2008] suggest that to assume that the correlation between \bar{v}_s and v_{st} can be estimated reasonably well there needs to be a minimum number of clusters (they cite 42). Most of the estimations proposed here have 51 clusters (states); an exception are those in Table 1.6, that have 14 clusters. It is not clear how to correct for the problem in this context, however. Angrist and Pischke [2008] suggest that at the minimum, one would want to show that the conclusions are consistent with the inferences that arise from group-level averages since this is a conservative approach. This is what I present in Column (1) of Table 1.6. The results of these group-level estimations are consistent with the micro-level estimates.

24. This section draws heavily from Angrist and Pischke [2008], pp.236-240

CHAPTER 2

Waves of Optimism: House Price Expectations and Credit Cycles

2.1 Introduction

The expectations of households and firms play a central role in macroeconomics. Ahead of the 2007 financial crisis American consumers channelled their savings into the real estate market largely because of the expectation of significantly positive returns on investment (Piazzesi and Schneider [2009], Case Shiller and Thompson [2012]). There is evidence that this optimistic attitude was shared by mortgage lenders: sophisticated investors appeared to be for the most part oblivious to the risk of a substantial downturn in the housing market (Coval et.al [2009], Foote et al. [2012]; Cheng et al. [2014], Chenenko et al. [2015]). The 2007-2008 crisis proved these expectations to be largely misguided.

Some models attempt to reconcile the burst of financial bubbles with rational expectations theory by framing them as the investors' reaction to *rare* events (Martin and Ventura [2011]; Caballero and Simsek [2013]).²⁵ This view is however hard to reconcile with the evidence that financial crises, and housing market crashes in particular, occur relatively frequently. Developed economies have experienced at least twenty housing market crashes in the post-World War II period.²⁶ Regional housing bubbles also have a long tradition: in the US, for example, they

²⁵. In other words, during boom phases agents are not blind to the possibility of a market downturn, but given the probability distribution of outcomes, it may be rational to invest in a given asset despite acknowledging its overvaluation. The burst of the bubble, on the other hand, occurs due to stochastic and exogenous processes, to which agents attach an extremely low probability ex-ante because they are rare events. Some classes of these models rely on frictions, others on asymmetric information: see Brunnermeier and Oehmke [2013] for a comprehensive review of the literature on rational bubbles.

²⁶. The IMF [2003] counts 20 of these episodes between 1970 and 2002, while Jordà et.al [2015] count 25 housing market crashes (for a similar cross-section of economies) between 1945 and 2013, all of which were followed by a recession.

date back all the way to the frontier land boom of the late 18th century (Glaeser [2013]).

The generalized underestimation of risk which occurred in the run-up to the 2007 financial crisis might therefore stem from some form of cognitive limitation: investors could be applying simple heuristics to predict price changes in the future (Glaeser [2013]). In particular, the excessive weight given to recent events when forming expectations might lead investors to highly discount the probability of a market downturn in good times, and *vice versa* (Gennaioli, Shleifer and Vishny [2015]; Bordalo Gennaioli and Shleifer [2016]). In this case, a form of “irrational exuberance” may have been a main driver of housing market dynamics in the pre-crisis period (Shiller [2015]).

This chapter studies how consumers form house price expectations, and whether expectations matter for their financial decisions. I provide three main contributions. First, I use the microdata contained in the Michigan Survey of Consumers and exploit its variation along the lines of geography and time to analyse how American households formed house price expectations between 2007 and 2014. I show that households have heterogeneous beliefs about the future of the housing market, which systematically depend upon household characteristics and upon the history of past house price realizations in the state of residence. Experiencing a state-level house price increase worth 1 percentage point (on average over four quarters) leads households to forecast a price increase of 0.6 percentage points, at the one-year horizon (10% of a standard deviation in the dependent variable). This coefficient falls in the range of the estimations provided by Case Shiller and Thompson [2012] and Kuchler and Zafar [2015] for a similar exercise, albeit they use different data.

Second, I show that extrapolation from recent house price changes induces a *systematic* bias in beliefs. I construct individual-level house price forecast errors and show that such errors are predictable from information publicly available at the time the forecast was made, in particular from *recent house price growth* in the state of residence. The estimated elasticity between recent state-level house price growth and changes in individual-level house price forecast errors (at the one-year horizon) is 0.67. In other words, if individuals experienced a state-level house price increase (decrease) in the past year worth 1 percentage point, their forecast errors about the future of the housing market tend to become 0.67 percentage points more optimistic (pessimistic). House price expectations seem therefore to follow a representativeness heuristic, as defined by Gennaioli Shleifer and Vishny

[2015], where agents overweight information they recently acquired when making predictions about the future.²⁷ The predictability of forecast errors from recent house price realizations is inconsistent with full information rational expectations theory, because in this framework expectations should be fully efficient at least with respect to past values of the variable being forecast (Muth [1961]; Lovell [1986]).

Considering the empirical relationship between leveraged housing bubbles and financial instability (Jordà et.al [2015a,b]) it is particularly interesting to study whether house price expectations directly affect mortgage leverage choices. Therefore, as a third contribution, I provide evidence that house price expectations help predict the mortgage credit cycle. I use the geographical information included in the Freddie Mac's Single Family Loan-Level Dataset, merged with state/quarter averages of house price expectations measured by the Michigan Survey, to test whether an increase in house price expectations leads to an increase in household leverage ratios.

The identification of a causal relationship between house price expectations and mortgage leverage is challenging, due to concerns over simultaneity and omitted variables. Therefore, in order to identify the effect of an exogenous shift in housing sentiment on the American mortgage market I exploit an instrumental variable strategy. Mian, Sufi and Khoshkhoh [2015] use the interaction of constituent ideology prior to presidential elections with election timing to show that more progressive (conservative) counties experience a positive shift of feelings about the government whenever the Democrats (Republicans) win the White House. In a similar fashion, I use the interaction of constituent ideology prior to the 2008 presidential election with election timing to show that more progressive states experienced a more positive *housing sentiment* shift around the time of the election, after controlling for pre-electoral trends. I show that this change in house price expectations can be considered exogenous to changes in fundamentals and to post-electoral policy changes that might affect the housing or mortgage markets directly. Most importantly, the shift in housing sentiment appears to be unaffected by changes in other sentiment variables, including feelings about the government as measured by Mian, Sufi and Khoshkhoh [2015].²⁸

27. Their definition follows Kahneman and Tversky [1972].

28. I borrow from Mian, Sufi and Khoshkhoh's [2015] methodology with the aim to answer a different question from them. They study how households' consumption responds to changes in feelings about the government; on the other this chapter is an analysis of how shifts in house price expectations affect the mortgage credit cycle.

I exploit this methodology to show that one standard deviation increase in state-level house price expectations (1.7 percentage points between 2007 and 2014) generates an increase in individual-level loan-to-value ratios worth 1 percentage point, or 6% of a standard deviation in the dependent variable. This result holds to controls for loan and borrower characteristics, aggregate time trends, state-level fixed effects and state-level house price growth. To my knowledge, this is the first contribution providing empirical evidence that house price expectations directly affect mortgage leverage ratios.

This chapter draws inspiration from several strands of literature. In particular, my work is closely related to the empirical efforts analysing the expectation formation process, which generally shows that expectations are heterogeneous among agents (Carroll [2003]; Branch [2004]; Souleles [2004]). Recent evidence strongly points in the direction of an extrapolative bias affecting different types of expectations (Malmendier and Nagel [2016]; Madeira and Zafar [2015]; Greenwood and Shleifer [2014]).²⁹ There are not many studies of house price expectations, however. Case Shiller and Thompson [2012] describe house price forecasts using proprietary data on four US metropolitan areas before the crisis, and find evidence of unrealistic five-year expectations. Bover [2015] focuses on the Spanish case, and shows how expectations are heterogeneous and depend upon household-specific characteristics. Kuchler and Zafar [2015] provide evidence that house price expectations depend on past housing returns.

I extend this literature by describing how house price expectations are formed using a publicly available data source which is representative of the American population. Moreover, by focusing on individual-level forecast errors, this chapter is to my knowledge the first to provide evidence of a *systematic* bias within house price expectations formed by American consumers.³⁰ This bias has been the object of speculation before (Case Shiller and Thompson [2012]; Shiller [2015]), but was never formally quantified.

My work is also related to the literature evaluating whether sentiment has any real effects on consumer and investor behaviour. Using the Michigan Survey of Consumers, Souleles [2004] shows that people's expectations are biased and

²⁹. Malmendier and Nagel [2011;2016] show that expectations and financial choices depend more strongly on lifetime experiences than on other publicly available data. Madeira and Zafar [2015] confirm this finding with respect to short-term inflation expectations and find that publicly available information matters more for longer horizons. However, evidence of an extrapolative bias is not confined only to inflation expectations. Greenwood and Shleifer [2014] show expected stock market returns are extrapolative in nature, and as such incompatible with rational expectations models of returns.

³⁰. The only other contributions testing the rationality of house price expectations is, to my knowledge, Zhang[2016], who focuses on professional forecasters. Zhang[2016] finds evidence of systematic over-optimistic forecasts, but not of inefficiency.

inefficient, and nevertheless sentiment helps forecasting consumption growth. De Nardi et al. [2011] find that in the context of a permanent income model, the fall in income and wealth expectations after the 2007 financial crisis can explain the post-recession consumption drop in its entirety. Several other studies have linked experiences, beliefs, and financial market decisions (Chernenko Hanson and Sunderam [2015] ; Giannetti and Wang [2014]; Malmendier and Nagel [2011; 2016]; Gennaioli Ma and Shleifer [2015]).³¹

However, empirical work on the macroeconomic effects of house price expectations is recent, mainly due to data limitations (as the collection of these surveys only began with the financial crisis).³² A recent applied literature analyses the feedback effects of house price expectations, or their capacity to be self-fulfilling prophecies. Lambertini et al. [2013] use a VAR approach to show that during housing booms, expectations about future house price growth account for a large fraction of macroeconomic fluctuations. Ling et.al [2015] show that changing sentiment from homebuyers, home builders, and lenders predicts house price appreciation in the following quarters. They also show that this feedback mechanism between sentiment and house prices can explain the persistence in house price movements along the boom and bust cycle. Soo [2013] develops an original measure of house price sentiment based on local area newspaper articles and finds that sentiment can predict a substantial fraction of the subsequent variation in house price growth. On a similar note, Wang [2014] finds that in states where homeowners overestimate the current market value of their home (a proxy for over-confidence), housing returns in the following year are higher. Overall, the evidence seems to point towards a strong feedback effect of house price sentiment on housing market equilibria. However, this paper is the first contribution to show that house price expectations have a direct effect also on the mortgage markets.

The chapter proceeds as follows. Section 2.2 describes the expectations data, and presents the results related to how house price expectations depend on individual-level characteristics. It also describes how I use the survey to analyse house price forecast errors. Sections 2.3 presents the mortgage-level data, the identification

31. Chernenko Hanson and Sunderam [2015] link personal investor experience to attitudes about investing in non-prime mortgages during the 2003-2007 credit boom. This evidence confirms the findings of Malmendier and Nagel [2016], who show that lifetime inflation experiences matter for the choice of mortgage products. Gennaioli Ma and Shleifer [2015] also show that past profitability is strongly correlated with CFOs' expectations about future profitability, and that that optimism in expectations affects their actual investment decisions. Giannetti and Wang [2014] find that experiencing corporate scandals reduces future individual participation in the stock market.

32. With the exception of Case Shiller and Thompson [2012] who focus on a few metropolitan areas using a proprietary dataset.

strategy, and the results on the effect of housing sentiment on mortgage leverage decisions. Section 2.4 concludes.

2.2 Empirical analysis of house price expectations

This section describes the Michigan Survey of Consumers, the data source used to analyse individual-level expectations. It also provides some descriptive analyses of the determinants of individual-level expectations and shows how the data can be used to test the rational expectations hypothesis. Finally, it presents the results of these tests.

2.2.1 Data: expectations in the Michigan Survey

Data on expectations comes from the University of Michigan Survey of Consumers, the source used to produce the Consumer Sentiment Index. This survey is nationally representative and has been conducted every month since 1978 on a rotating panel of about 6000 US households (500 per month).

The interviews are conducted with one individual per household and include household-level demographics such as income, educational attainment, and family composition, as well as a vast array of sentiment and expectations indicators. In particular, respondents are required to indicate their forecast of the one-year-ahead percentage change in inflation, personal income, and local area house prices.

These questions are phrased as:

“By about what percent do you expect prices of homes like yours in your community to go (up/down), on the average, over the next 12 months?”

Similar questions are asked about the development of personal income and inflation.³³ The descriptive statistics for this sample are provided in panel A of Table 2.1.

To analyse how individuals’ experiences and characteristics influence the expectation-formation process, I estimate the following equation:

$$Expectation_{ist} = a + \beta_1 \Gamma_{ist}^I + \beta_2 \Theta_{st}^I + \beta_3 \phi_s + \beta_4 \chi_t + \varepsilon_{ist} \quad (3)$$

³³ House price forecasts are only available since 2007 and only for homeowners. For consistency, I therefore drop non-homeowners from the sample altogether. This exclusion is feasible because homeowners are the majority of the survey sample (they constitute 78% of surveyed households).

Where the outcome variable is the individual-level expectation about the change in income, inflation, and house prices in 12 months for individual i living in state s during quarter t . Γ^I is a vector of individual characteristics, such as income, a variety of demographics, and recent experiences. Θ_{st}^I measures aggregate dynamics at the state level in a given quarter t , such as recent house price changes, or unemployment rates.³⁴ Quarter fixed-effects have the purpose of controlling for aggregate shocks affecting all states at the same time, and state fixed effects control for time-invariant factors that might affect all families living in the same state across time.

2.2.2 Results: determinants of individual expectations

Aggregate expectations on the growth rates of income, house prices, and inflation display a strong correlation with the US business cycle (Figure 2.1). All three indicators of consumer confidence drop in the aftermath of the 2007/2008 financial crisis. Income growth expectations drop from 3% per year in 2007 to about 1%, and start recovering only in 2013. House price growth expectations follow a similar pattern: they become negative in 2008 and stay negative until 2012. Throughout this time, American consumers were consistently expecting a wealth loss. Expectations about inflation rates, by comparison, have been remarkably stable throughout this time frame, slightly falling in 2009/2010 only to recover swiftly around 2011 and they have been averaging around a yearly 4% ever since.

Table 2.2 sheds some light on how household-level demographics are correlated with different measures of expectations. Richer and older couples have on average lower income expectations than their younger, poorer, and single individuals (Column 1). This probably reflects the lifecycle of earnings. On the other hand, men and people with a college degree expect their earnings to grow more than other demographic groups. People who report experiencing negative income shock in the previous year (measured as job loss or reduced wages/working hours) expect their income to grow 2.5 percentage points less than others. This is coherent with recent evidence showing that negative shocks at the personal level cast a shadow of pessimism on agents' beliefs about the future. For example, individuals who experienced a negative stock market shock are more risk-averse and less likely to predict high returns on investment (Malmendier and Nagel [2011]).

³⁴. Details of the state-level control variables can be found in Appendix A2.1 and the relative descriptive statistics in Panel C of Table 2.1.

The effect of unemployment rates confirms this intuition: one standard deviation increase in state-level unemployment rates (2%) reduces individual income expectations by 8 percentage points (57% of a standard deviation in the dependent variable). This can be considered evidence corroborating the findings of Kuchler and Zafar [2015], who find that experiencing unemployment systematically makes people more pessimistic about the future of the labour market.

Inflation expectations display different correlations with household demographics (Column 2). Richer and more educated males expect future inflation to be lower than poorer and less educated women, or older people. This is consistent with the results presented by Madeira and Zafar [2015], who find that women, ethnic minorities, less educated and lower-income people predict higher inflation, on average. They also find that these social groups are slower in updating their expectations, and make more prediction errors. Madeira and Zafar [2015] interpret their results as indicative of differentials in the ability to collect and process public information across different types of agents.

My results also indicate that stock owners expect lower levels of future inflation, which may be consistent with a theory of the heterogeneity in expectations being based on information. Stock market exposure may induce people to follow the financial news more closely, and this may in turn develop their ability to better assess market conditions. On the other hand, people who are more financially literate probably also self-select in stock ownership. Access to information, as well as information processing ability (financial literacy) may therefore play a crucial role in explaining heterogeneity in inflation expectations, as suggested by Burke and Manz [2014]. People who recently experienced a negative income shock, on the other hand, forecast future inflation to be higher, giving further credit to the idea that personal experiences matter for relative optimism/pessimism about the future.

Column 3 of Table 2.2 analyses how house price expectation depend upon household demographics and past experiences. Richer households, men, and college graduates expect house prices to grow more, as do people who own stocks. This might in part be due to unobserved within-state heterogeneity: these households may be more likely to reside in metropolitan geographical areas, where house prices are likely to grow more than the state average, all else held equal.³⁵ On the other hand, it may also be that these households are actually better informed, and correctly anticipated a more rapid house price growth in the post-crisis period.

35. I observe the state of residence for any given household, but not the county or ZIP code.

Less informed households may be more prone to cognitive biases and may be slower in updating expectations, as suggested by Madeira and Zafar [2015]. They may therefore have projected the housing market shock to continue well beyond the crisis of 2007-2011.³⁶ Once again, people who recently experienced a negative income shock are less optimistic about the future, in housing markets as well as income and inflation.

An interesting result of this specification is that the average yearly house price growth in the state of residence (measured in the quarter preceding the interview) is a strong predictor of expectations about future house price growth. A household experiencing a 1 percentage point increase in state-level house prices in the previous four quarters predicts the one-year-ahead increase in local house prices to be 0.5 percentage points higher, or about 10% of a standard deviation in the dependent variable (significant at the 1% level). This result lies in between the elasticity of 0.23 estimated by Case Shiller and Thompson [2012] and that of 0.9 estimated by Kuchler and Zafar [2015].³⁷

This evidence suggests an extrapolative pattern in house price expectations: if individuals experience house price growth in their state of residence, they expect the trend to continue in the near future. These results are consistent with other recent empirical studies focusing on different kinds of expectations and provide evidence for an extrapolative component of investors' beliefs about the future that largely depends on recent experiences (Malmendier and Nagel [2016]; Madeira and Zafar [2015]; Greenwood and Shleifer [2015]; Kuchler and Zafar [2015]).

Overall, this section shows that expectations are highly heterogeneous across households. Different demographic groups display systematic differences in the way they think about the future. This seems to contradict the tenet that private information plays no role in the expectation formation process, and that therefore all expectations can be approximated by those of a representative agent (Muth [1961]). On the other hand, it does not necessarily contradict the hypothesis that expectations are formed efficiently overall, since from the point of view of the individual it may be optimal to choose different forecasting methods depending on personal circumstances, or different individuals might have differential access to information (Pesaran and Weale [2006]).

36. In other words, they could have behaved according to the representativeness heuristic, described by Gennaioli Shleifer and Vishny [2015]. In this framework, Bayesian agents biased by representativeness only react to a *series* of good/bad news and overweight recent trends (as opposed to considering all historical information) when making predictions about the future.

37. Kuchler and Zafar [2015] use a different dataset, study a different time frame (2012 onwards), and measure house prices at the local (ZIP) code area rather than at the state level. Case Shiller and Thompson [2012] focus on the pre-crisis period and on four metropolitan areas.

2.2.3 Testing the rationality of expectations: methodology

Table 2.2 shows that people form expectations about the future based on the information available to them at the time they make the forecast. Economic theory adds to this tenet the notion of *optimality* in the use of publicly available information: that is to say, individuals might make mistakes in their predictions, but the economic system in the aggregate does not waste information. In this sense, expectations are assumed to be rational, or consistent with the predictions of the relevant economic theory (Muth [1961]).

Muth's [1961] original theory postulates that private information plays no role in the formation of macroeconomic expectations. Moreover, expectations should be *fully efficient* with respect to publicly available information. Given a variable Y , its value at time t (Y_t) should be perfectly predicted by the ex-ante expectations of the representative agent, defined as $E_{t-n}(Y_t)$. Any vector of public information available to the agent at time $t-n$ (X_{t-n}) should have no additional explanatory power towards Y_t . Formally:

$$Y_t = \alpha + \beta_1 E_{t-n}(Y_t) + \beta_2 X_{t-n} + \varepsilon_t \quad (4)$$

with $\alpha = \beta_2 = 0$; $\beta_1 = 1$; $E(\varepsilon_t) = 0$.

Forecasts may diverge from realisations, but the errors will average out to zero over time, and they won't be *systematic*. This, in turn, implies orthogonality between ex-post forecast errors (FE_t) and all public information available to the agent at the time the forecast $E_{t-n}(Y_t)$ was made:

$$FE_t = Y_t - E_{t-n}(Y_t) = \alpha + \beta_2 X_{t-n} + \varepsilon_t \quad (5)$$

with $\alpha = \beta_1 = 0$; $E(\varepsilon_t) = 0$.

In other words, under rational expectations forecast errors should be unpredictable given the set of public information available to the agent at the time the prediction was made (Muth [1961]; Lovell [1986]).

In order to test whether house price expectations are formed in a way that is consistent with the full-information rational expectations model, I construct individual-level house price *forecast errors*, defined as follows:

$$FE_{ist} = E_{ist-4}(HPI_{st}) - HPI_{st} \quad (6)$$

Where $E_{ist-4}(HPI_{st})$ is the expectation that individual i living in state s at quarter $t-4$ has about house price growth in state s at time t (percentage house price growth in one year). This forecast is compared with the actual annualized change in house prices recorded in quarter t for state s , as measured by the Federal Housing Finance Agency (FHFA) quarterly repeated sales house price index (HPI_{st}).³⁸ FE_{ist} therefore represents individual-level *forecast errors*: unlike in equation (5), a larger value implies over-optimism.

Note that equation (6) introduces individual-level heterogeneity in the definition of forecast errors, which was absent from equation (5). The presence of individual-level heterogeneity in expectations, described in the previous section of this chapter, suggests that also the forecast errors (FE_{ist}) are unlikely to be orthogonal to the private information set, defined by individual characteristics. It is not clear yet how to test for rationality in the presence of individual-level heterogeneity (Pesaran and Weale [2006]). Heterogeneous individuals may have different information processing costs, and it may be optimal for them to choose different forecasting methods (Pesaran and Weale [2006]). Moreover, agents may have differential access to information. The definition of rational expectations in the context of heterogeneous information and information processing capacity constitutes a very interesting avenue of research. However, it is beyond the scope of this chapter to analyse this matter in great detail. Instead, I will focus on whether *public* information (specifically past house price growth in the area of residence) is processed efficiently on average. I therefore exploit the panel component of the survey, which provides two observations per individual, to study how public information translates into changes in individual-level forecast errors.³⁹ To do so, I use a model in first-differences:

$$\Delta FE_{ist} = a + \beta_1 \Delta X_{ist}^I + \beta_2 \Delta \Theta_{st-n}^I + \gamma_t + \varphi_s + v_{ist} \quad (7)$$

Where $\Delta FE_{ist} = FE_{ist} - FE_{ist-2} = \{[E_{ist-4}(HPI_{st}) - HPI_{st}] - [E_{ist-6}(HPI_{st-2}) - HPI_{st-2}]\}$ is the difference between individual i 's forecast errors between the first and the second interview (which are two quarters apart from each other) and s and t indicate state and quarter, respectively. ΔX_{ist}^I is a vector of changes in family-specific controls between the first interview and the second one, such as household income,

³⁸. Information about the house price index, together with other aggregate controls, can be found in Appendix A2.1.

³⁹. An alternative would be to estimate a model with state-level averages. The results are very similar in nature, magnitude and significance, with respect to this model in changes. I therefore prefer to maintain micro-level variation and use a model in changes at the individual level instead.

plus the same household-level demographics used in Equation (3) measured at the time of the latest interview. $\Delta\Theta^I_{st-n}$ defines changes in state/quarter variables, such as the average yearly growth in house price (measured in the quarter prior to each interview).

By first-differencing the outcome variable, the model controls for all time invariant household-level characteristics related to idiosyncratic perceptions of the housing market.⁴⁰ First-differencing should also control for all individual level heterogeneity that can be reasonably assumed to be constant for a given individual within six months, such as information processing capacity and financial literacy.

The model in Equation (7) allows for the identification of systematic components in consumers' forecast errors. Significant coefficients on any variable within the information set available to the agent at the time they produce the forecast (any β_1 or $\beta_2 \neq 0$), imply a departure from a strong form of rational expectations (Lovell [1986]). On the other hand, for even a weak form of rational expectations to hold, the prediction errors must at least be independent from historical information on prior realisations of the variable being forecast: formally, β_2 must be equal to zero, whenever Θ_{st-n} measures past house price realisations (Lovell [1986]).⁴¹

2.2.4 Results: house price forecast errors

Figure 2.2 shows that individual-level forecast errors do not cancel each other out in the aggregate, since they are significantly different from zero in each quarter. Also, they display a strong time component in this sample: they are systematically positive, implying excessive optimism during the crisis and the recession (2007q1 until 2010q4) and consistently negative afterwards, indicating that American consumers have on average underestimated the recent recovery of the US housing market.

40. For example, individuals might be forecasting house price growth for their local area of residence (ZIP code or city) rather than for their state. Since I construct individual-level forecast errors as the difference between the individual-level expectation and the state-level realization, the dependent variable might contain measurement error (the difference between local and state-level house price growth). As long as this difference is constant over six months, the model in changes should take into account this unobserved variation. However, even if the difference between state and local area house price growth were to change over time, this difference will appear as measurement error in the *dependent* variable. As long as the measurement error in the dependent variable is uncorrelated with the right-hand side of Equation (7), the estimation will be consistent.

41. This version of the rational expectations hypothesis is weak in that it only requires the agent to efficiently process the information related to the historical realizations of the variable s/he is forecasting rather than *all* available public information.

The first two columns of Table 2.3 analyses how forecast errors depend on household characteristics. Column (1) describes forecast precision: the dependent variable is the absolute value of forecast errors. The closer this value is to zero, the higher the precision of the forecast. A larger value therefore implies higher inaccuracy. Richer households, men, and households who invest in the stock market have more accurate estimates about the future of the housing market. The effect of owning stocks is small, but strongly significant: stock owners make predictions that are on average 0.3 percentage points more accurate than non-stock owners (or 5.7% of a standard deviation in the dependent variable). This is also true of more educated families: the effect of having a college degree improves the accuracy of the house price forecast by 0.14 percentage points, or 2.6% of a standard deviation in the dependent variable. This seems consistent with what Madeira and Zafar [2015] find about household-level heterogeneity in inflation expectations: women, less educated people, and poorer households tend to have more imprecise forecasts. This result provides further support for the hypothesis that access to information, or the ability to process it, might play a crucial role in the expectation formation process. People who recently experienced negative income shocks also tend to have less precise forecasts.

However, forecast errors can also be analysed with respect to their relative degree of optimism and pessimism, rather than in absolute values. Column (2) shows the results of a model with the forecast error defined as in Equation (6): a positive value in the dependent variable now implies excessive optimism about the future of local house prices. Richer households, men, and college graduates tend to have more positive forecast errors: in other words they are wrong less often (as shown in Column 1), but when they are, their mistakes are on the optimistic side. A negative income shock, on the other hand, makes people excessively pessimistic about housing market returns (Column 2): a family declaring a negative income shock in the previous year predicts a house price growth at the on-year horizon 0.74 percentage points lower than the actual realization (about 10% of a standard deviation in the dependent variable).

The predictability of forecast errors from household-level demographics should however not necessarily be interpreted as evidence against the rational expectations hypothesis. It might be that different individuals have differential access to information; or that private information plays a role in determining house price expectations; or, again, that the same public information is optimally processed in different ways by different individuals. So far, consensus has not emerged yet

on how to distinguish between these alternative hypotheses (Pesaran and Weale [2006]).

Column 3 therefore estimates a model in first-differences in order to control for time-invariant individual-level heterogeneity, in an attempt to mitigate any differences due to *private* information and household-specific characteristics. This specification attempts to evaluate whether the economy in the aggregate processes *public* information efficiently, by analysing whether forecast errors are efficient with respect to past information about local area house price growth.

Column (3) shows that a recent history of housing appreciation (average house price growth in the previous four quarters) is indeed a strong predictor of the change in individual-level house price forecast errors.⁴² An increase in state-level house prices worth 1 percentage point in the year before the forecast was initially made is correlated with an increase in individual forecast errors worth 0.67 percentage points. In other words, if individuals experienced a state-level house price increase (decrease) in the past year worth 1 percentage point, their forecast errors about the future of the housing market tend to become 0.67 percentage points more optimistic (pessimistic). A similar increase in personal income affects forecast errors by 0.4 percentage points. These results hold when controlling for time trends and state-specific characteristics, and individual-household fixed effects. The inclusion of quarter fixed effects should also rule out the possibility that the forecast errors may be due to unexpected macroeconomic shocks, since all aggregate time trends are taken into account. It is still possible that these results reflect unobserved state/quarter shocks, but considering that these are short-term expectations, the possibility that a repeated series of unanticipated shocks at the state/quarter level is driving the result seems unlikely.

The latter result suggests that households may not be efficiently processing public information about house price growth. People's forecast errors have a tendency to become over-optimistic when they experience house price growth, and *vice versa*. In other words, they attach too much weight to recent house price movements. This is only a baseline result, however it could be extended in several directions by exploiting both time and geographic heterogeneity in forecasts. For example, it would be interesting to evaluate whether the bias appears to be stronger around the financial crisis, or whether people extrapolate from losses as much as from gains. Another dimension of heterogeneity could be geography, since the extrap-

42. To avoid simultaneity, past housing appreciation is measured as the average yearly house price growth measured in the quarter ahead of the interview.

olative bias could depend upon certain characteristics of a state: for example the extent to which the state has been subject to major housing market crashes in the past, or its history of house prices volatility. Finally, it would be interesting to study how individual characteristics interact with the generalized tendency towards an extrapolative bias.

However, at the very least, the result in Column 3 suggests that consumers' house price expectations are not formed efficiently and are subject to some form of cognitive limitation. In particular, house price expectations seem to violate even a weak form of the rational expectations hypothesis, since they are not efficient with respect to past realizations of the variable being forecast. People expect recent price movements to continue the future, leading to systematic errors. This result, combined with the evidence of heterogeneity in expectations and forecast errors, casts a shadow of doubt on the applicability of full-information rational expectations theory to housing markets. Taken literally, this result instead supports the hypothesis that this asset market may be subject to purely belief-driven boom and bust cycles, in which prices can be largely detached from fundamentals and be subject to excessive volatility (Gennaioli, Shleifer and Vishny [2015]).

2.3 House price expectations and the credit cycle

The expectation formation process matters because beliefs might affect the business cycle. Souleles [2004] provides an analysis of how income and inflation expectations in the Michigan Survey translate into households' propensity to consume. However, given the link between housing collateral and mortgage debt (Mian and Sufi [2011]), it might be particularly interesting to study whether house price expectations affect mortgage choices. In this section, I first describe the problem of identifying a causal relationship between house price expectations and credit markets, and present the empirical strategy I will use to address this problem. I then describe the mortgage data, which is derived from a publicly available, lender-side source. Finally, I present evidence of the empirical relationship between house price expectations and mortgage borrowing.

2.3.1 The identification problem: IV strategy

Mortgage leverage played a prominent role ahead of the financial crisis and its aftermath: it increased the likelihood of default on mortgages (Mian and Sufi [2009]) and was one of the main drivers of the recession (Mian and Sufi [2011]).

House price expectations display a strong positive correlation over time with average mortgage leverage recorded among American households (Figure 2.3). If individuals expect the value of their properties to rise, they might borrow against the expected increase in home equity, because a part of the loan will be automatically repaid by the price increase. At the same time banks might be willing to lend larger sums, because of the expectations of higher collateral in the near future.

The Michigan Survey of Consumers does not provide data on financial liabilities, but information on mortgage leverage at the household level is available from a variety of public sources. This mortgage-level data can be merged with state/quarter averages of house price expectations observed in the Michigan Survey of Consumer to estimate a model of the type:

$$LTV_{ist} = a + \beta_1 Exp_{st} + \beta_2 \Lambda_{ist}^I + \beta_3 X_{st}^I + \varphi_s + \chi_t + v_{ist} \quad (8)$$

Where LTV_{ist} is the individual mortgage loan-to-value ratio, a measure of leverage, for family i residing in state s in quarter t . Exp_{st} defines the weighted average of expectations in state s at quarter t recorded by the Michigan Survey of Consumers. Λ_{ist}^I is a vector of household-level controls, which includes the credit score of the borrower, interest rate, length and purpose of the loan. X_{st}^I defines control variables recorded at the state/quarter cell, which might contemporaneously affect expectations and the dependent variable (such as recent state-level house price growth).

Year fixed effects allow to control for economy-wide shocks, for example federal policy changes affecting all states at the same time. Geographical area fixed-effects instead control for time-invariant state-specific characteristics, which could be correlated with both sentiment and mortgage markets. So the coefficient β_1 measures the effect of a within-state change in expectations over time, washed out all variation due to unobserved state-specific characteristics common to all households living in state s and all aggregate shocks occurring in the economy at time t .

However, Equation (8) is subject to several endogeneity concerns. First of all, there is a reverse causality problem: a higher availability of credit is likely to trigger a change in aggregate expectations about the future of house prices. Furthermore, expectations and outcomes are likely to be simultaneously affected by unobserved fundamentals. In order to identify the effect of a shift in house price expectations, I therefore rely on an instrumental variable strategy. Mian,

Sufi and Khoshkhrou [2015] show that that the ideological predisposition of residents in a county (Republican VS Democrats) is a strong predictor of within-county changes in sentiments regarding government policy, anytime there is a change of government in the White House. In particular, Republican-leaning counties become more pessimistic about government policy when Democrats win the presidential elections, and *vice versa*. However there is evidence that large-scale electoral events shift all measures of consumer sentiment, and not just views of the government. Gerbert and Huber [2009; 2010] exploit an unanticipated change in political power (the Democrat takeover of Congress occurring in 2006) to show that pre-electoral political leanings have a strong effect on the changes in *economic opinions* after the election. Immediately after the event, Democrats become more optimistic about the general economy than they were the month before, and Republicans' sentiment shifts in the opposite direction.

This suggests that housing expectations might also be subject to changes around election time. The housing market could be particularly affected by elections due to the role of pre-electoral uncertainty. Pre-electoral uncertainty may reduce investments that are costly to reverse: Canes-Wrone and Park [2014] find evidence of this effect across the US at the turn of the 2008 Presidential election. The extent of the reversal of this uncertainty after the election may depend upon the ex-ante political views of a certain electorate and interact with the party change at the White House. I exploit this idea to evaluate how the change in house price expectations following the 2008 presidential election affects mortgage leverage ratios.

My empirical strategy is formally expressed by equations (9) and (10).⁴³ The first stage relationship measures the within-state change in expectations occurring after the presidential election and takes the form:

$$E_{st} = \beta_1 Z_{st} + \beta_2 X_{st}^I + \beta_3 \Lambda_{ist}^I + \beta_4 \phi_s + \beta_5 \chi_t + \varepsilon_{st} \quad (9)$$

Where E_{st} are state-level expectations about one-year change in house prices, measured as a weighted average of the individual level forecasts provided by the Michigan Survey in a given state s and quarter t . $Z_{st} = D_s * \delta_t$ is the interaction term between the vote share for the Democratic Party in a given presidential election (D_s) and the post electoral period δ_t (the quarter of the election is excluded from the analysis). State fixed effects, ϕ_s , capture time-invariant state charac-

43. I apply a very similar methodology to Mian, Sufi and Khoshkhrou [2015], although I rely on state-level measures of sentiment and political leanings, rather than on county-level data. Also, my time series is shorter, because the Michigan Survey only began collecting information on house price expectations in 2007.

teristics (including the voting share for the Democratic party in 2008, D_s) while quarter fixed effects, χ_t , control for economy-wide time trends such as the US-wide shift in housing sentiment occurring in 2007 (and take into account the time varying term δ_t). If Z_{st} is significant and positive, progressive states get more optimistic about the housing market than conservative states, and this shift occurs *after* the electoral period. It is important to stress that the vote shares for the Democratic party are *not* assumed exogenous in this model: partisan leanings can be strongly correlated with long-term housing price dynamics, such as the willingness to issue new building permits (Kahn [2011]). The validity of the instrumental variable strategy relies on the exogeneity of the *interaction* between partisanship and electoral timing. The vector X_{st}^I is a set of variables that proxy for changes in fundamentals which could impact states exactly at the time of the elections (for example changes in house prices). The vector Λ_{ist}^I includes individual-level characteristics of borrowers in each state: income, age, credit score of the borrower, and some characteristics of the loan (length in years, interest rate, type and purpose of the mortgage, use of the property).⁴⁴ The inclusion of these variables has the purpose of building further credibility to the orthogonality condition: the exclusion restriction is valid after partialling out for these shocks in fundamentals.

I will run a series of robustness tests on the first stage relationship to show that the switch in housing sentiment can be considered exogenous to a set of macroeconomic fundamentals. I also show that the switch cannot be attributed to changes in state-level housing policy taking place after the elections and that it is robust to shifts in other expectations and sentiment variables occurring at the same time, including views of the government (as measured by Mian, Sufi and Khoshkhoh [2015]).

The second stage relationship exploits mortgage-level data to measure how house price expectations affect borrowing/lending. This relationship is defined as follows:

$$LTV_{ist} = a + \beta_1 \widehat{E}_{st} + \beta_2 \Lambda_{ist}^I + \beta_3 X_{st}^I + \beta_4 \varphi_s + \beta_5 \chi_t + v_{ist} \quad (10)$$

Where LTV_{ist} is loan-to-value ratio for household i , in state s , at time t . The dependent variable is regressed upon the same set of controls in (9), with the housing sentiment variable at time t (\widehat{E}_{st}) instrumented by Z_{st} . The second stage of this instrumental variable strategy deserves however some further discussion.

44. The inclusion of this vector in the first-stage relationship is necessary for the consistency of the IV estimator.

Mian, Sufi and Khoshkhoh [2015] use the change in presidency to evaluate how sentiment *towards the government* affects household consumption. This might undermine the credibility of my empirical results, because the exclusion restriction in Equation (10) might be violated. In particular, the concern might be that views of government are driving the second stage results, rather than house price expectations.

While I recognize the limitations associated with this particular instrument, it is hard to isolate an alternative variable that only affects sentiment and does not impact the mortgage and housing markets directly. Moreover, I believe this particular concern about the validity of the instrument to be of second-order relevance for two main reasons. The first is that other sentiment variables, including views about the government, display a much lower correlation with household loan-to-value ratios than housing sentiment does. The second (and perhaps more convincing) reason is the fact that Mian, Sufi and Khoshkhoh [2015] show that shifts in government sentiment after the election have *no significant effect* on household consumption over the same time period. If views of the government do not shift short-term consumption habits, there is no reason to believe that they will influence long-term saving decisions such as mortgage borrowing, and drive my second-stage results.

2.3.2 Mortgage-level data

The Michigan Survey does not include individual-level data on households' balance sheets. Therefore, in order to identify whether house price expectations matter for mortgage leverage choices I have to rely on a different data source.

The data on individual-level mortgage transactions comes from the Single Family Loan-Level Dataset, provided by the Federal Home Loan Mortgage Corporation (Freddie Mac). While some surveys collect information about American households' financial liabilities, using mortgage information provided by the lender provides several advantages. The first is coverage: Freddie Mac's database collects information about over 20 million residential mortgages securitized across the United States between 1999 and 2015. Freddie Mac's share of mortgage-backed securities currently corresponds to roughly a third of the American market in terms of number of loans, and 14% in terms of volume.⁴⁵ The second advantage of this dataset is its precision and quality. As this is lender-level data, it is much less

⁴⁵. <http://www.freddiemac.com/investors/pdffiles/investor-presentation.pdf>; Federal Reserve Board Data, Mortgage Debt Outstanding, March 2016.

likely to contain measurement error. It also provides information unavailable in most surveys, such as credit score of the borrower or the length of the mortgage in years.

Freddie Mac provides a sample of about fifty-thousand observations per year which are randomly drawn from the overall population. I rely on this sample because it makes the estimations less computationally intensive while matching the moments of the distribution of the overall population very closely.⁴⁶

After the Federal Housing Finance Regulatory Reform Act of 2008, the Agencies (as Freddie Mac and Fannie Mae are commonly referred to) were put under federal administration and are now running under the conservatorship of the Federal Housing Finance Agency (FHFA). Since the federal government is ultimately responsible for the Agencies' solvency, both have strict rules about the characteristics of the mortgages that fall under their umbrella. Loan values cannot exceed certain nominal limits, which are determined annually by the FHFA, depending on the geographical area where the house is located. The Agencies are also required to back only "prime" mortgages, and jumbo loans are excluded from their portfolios. This dataset in particular is composed only of 30-year fixed-rate single-family mortgages, which nevertheless constitute the most common type of mortgage on the American market, making up on average 83% of the stock of loans originated in a given year (Fuster and Vickery [2015]). In this sense, this dataset represents the most conservative side of the American mortgage market, both in terms of lending risk and overall leverage.

The descriptive statistics for this sample can be found in panel B of Table (2.1). The outcome variable I will consider is the individual mortgage loan-to-value ratio. This is the ratio of the loan to the value of the property as appraised by the lender (or the original property value at the time of purchase, if the owner can prove that the value of the property has not declined since then). Clearly, this is an equilibrium variable, because it jointly reflects credit supply and credit demand. The average mortgage securitized by Freddie Mac in this time frame is worth 69% of the property value and its length is 26 years. About 47% of the families have not been homeowners in the past three years and as such are labelled in the database as first-time home buyers. The vast majority of borrowers live in the property, since investment loans are only 6% of the total. On the other hand, a large fraction of the loans have the purpose of refinancing since purchase mortgages are the minority (37%).

46. All relevant comparisons between the sample and the population are included in the material provided by Freddie Mac together with the dataset.

2.3.3 First stage: housing sentiment and elections

Column 1 of Table 2.4 estimates Equation (9), the first-stage relationship. The relationship between house price expectations, partisan leanings and electoral outcomes is strong: the interaction between ex-ante state-level voting share for the Democratic Party, and a dummy indicating the post 2008q4 period, displays an elasticity of +0.10 (significant at 1 percent level) after controlling for pre-electoral trends.

Column 2 controls for whether this shock was due to changes in fundamentals. For example, if Democratic states experienced an income shock at the time of the elections, or a stronger house price growth, Column 1 would be capturing a spurious correlation between the electoral outcomes and the sentiment variable. As in Table 2.2, price expectations are strongly correlated with past house price growth, but other controls, such as changes in aggregate income, unemployment rates are not significant. The coefficient on electoral outcomes is still positive and significant at the 1% level, even if its magnitude is smaller than the one estimated in Column 1 (+0.7). Time trends are taken into account, not only by a dummy indicating the post electoral period (after the third quarter of 2008) but also by quarter fixed-effects. State fixed-effects are included in order to control for time-invariant characteristics of a state, such as geographical constraints that might limit housing supply elasticity (Saiz [2010]).

Given that the shift in presidency in 2008 affects house price expectations, if house price expectations affect mortgages a relationship between the potential instrument and loan-to-value ratios should emerge. Column 3 provides the estimation of this reduced-form equation. The coefficient on the interaction term is positive and significant (+0.05) at 5% level. The Democratic vote share in a given state is positively associated with an increase in the individual-level loan-to-value ratios, in the post-electoral period. In the remaining part of this section, I run a series of robustness tests aimed at verifying the validity of this instrument and testing the credibility of the exclusion restriction.

Post-electoral policy changes

A cause of concern with the first stage estimation presented in Table 2.4 is that local political leanings might directly affect housing markets in the post-electoral period, for example by changing the local housing policy. Local housing policy might lead people to act differently with respect to the housing market, regardless of house price expectations. In this case the exclusion restriction would be violated.

Table 2.5 runs some robustness checks, testing whether Democratic-leaning states experience relevant policy changes in the housing sector in the post-electoral period. Columns 1 and 2 estimate whether more progressive states experienced a change in federal or local public transfers devoted to housing after 2008. For example, if Democratic states changed the provision of housing benefits after the 2008 election, a fraction of poorer citizens might have been pushed into the private residential market, changing the overall leverage ratios. The evidence in Column 1, however, suggests that Democratic states did not change housing benefits in the post-electoral period. The percentage of citizens relying on public housing is not significantly affected by political leanings (Column 1), nor is the number of people relying on rent subsidies (Column 2).

A different kind of concern relates to Democratic states receiving a more favorable treatment in the post electoral period from the Federal government, for example in terms of real estate taxation. If Democratic constituencies experienced a decrease in property taxes after the election, this might increase people's willingness to buy a house, and possibly increase average loan-to-value ratios. Column 3 shows that this is not the case: the coefficient of the interaction term on average property taxes (in logs) is positive and not significant.

Finally, political views might change the housing supply dynamic after elections such as a change in the regulatory framework. Column 4 shows that the housing regulatory framework, proxied by the number of building permits issued in a given state/year, is not significantly affected by the interaction between Democratic vote shares and the post-electoral period. Ex-ante political leaning doesn't seem to change housing supply dynamics also when looking at construction costs (Column 5): the average wages in the construction sector, which are the component of building costs more likely to differ across states, are also unaffected by partisanship and election timing.

Overall, there is no clear relationship between ex-ante political views and policy changes in the post-electoral period which might impact the housing or mortgage markets directly.

Other sentiment variables

Table 2.4 controls for some of the fundamentals which might affect the regional economies at the time of the elections. However, the change in party at the White House might affect other sentiment variables, which could also have an impact on borrowing and saving decisions. It is important to understand if the post-

electoral shift in housing sentiment is not actually concealing a more general optimism about the economy. In particular, Mian, Sufi and Khoshkhoh [2015] show that the percentage of the population expressing a positive opinion about the government shifts dramatically after the US presidential elections. Other expectations about macroeconomic policy, such as inflation rates and interest rates, might also change substantially with a shift in government. Therefore, it is necessary to clarify the role of other expectations in the first-stage relationship.

Table 2.6 addresses this issue by studying the state-quarter average of other sentiment variables measured by the Michigan Survey of Consumers. Column 1 shows that the interaction between ex-ante political views and electoral outcomes has no significant effect on personal income expectations. The coefficient is positive, but not significant. However, more progressive states view the election of a Democratic president as beneficial for their income in real terms, as they expect the inflation rate to be lower in the subsequent 12 months (Column 2). Interestingly, there is no expectation of a change in interest rates following the election: Column 3 shows that the coefficient associated with the interaction term is not significant. This is reassuring, as interest rate expectations may be a main driver of credit market dynamics. Consistently with Mian, Sufi and Khoshkhoh [2015], I find that the proportion of individuals reporting a positive view of the government increases substantially after 2008, depending on the share of votes that the Democratic party received in the election (Column 4).

In sum, partisan leanings are a strong predictor of changes in inflation expectations and views of the government, after the 2008 presidential election. However, the first-stage relationship between elections and housing sentiment holds after controlling for these factors: Column 5 of Table 2.6 shows that the effect of the interaction term on house price expectations is still positive (+0.04) and strongly statistically significant (1% level) after controlling for other sentiment variables and for views about government policy.

Placebo test

In order for the instrument to be valid it is necessary to provide a convincing case that this observed shift in house price expectations is due to a pure sentiment shift, rather than a change in other variables. To this purpose, table 2.7 provides some placebo tests, checking whether the voting share for the Democratic party predicts sentiment shifts in times other than the 2008 election. Column 1 presents the baseline result (the within-state change in housing sentiment occurring after the 2008 presidential election as a function of the voting shares for the Democratic

party in the same state). The coefficient is positive (+0.07) and statistically significant at the 1% level. To test whether Democratic-leaning states experienced other sentiment shifts in the years after the 2008 election, I regress the within-state change in house price expectations on the interaction between voting shares for the Democratic party in 2008 and dummy variables indicating the years post 2009 and post 2010, respectively. In other words, while Column 1 observes the within-state change comparing the quarters before 2008 with all the quarters after, Columns 2 and 3 only compare the change occurring between 2009-2014 and 2010-2014, respectively. Democratic leaning states experienced a within-state sentiment shift (with respect to pre-2009 trends) also in late 2009, but its magnitude (+0.02) is less than a third of what they experienced in late 2008, and the estimate is quite imprecise. Also, Democratic states do not experience any significant sentiment shift in late 2010, with respect to pre-2010 trends. This implies that the change in house price expectations occurring in Democratic states in the post-electoral period was taking place mostly around the 2008 election.

Column 4 provides an additional robustness test, using the 2012 presidential election. As Mian, Sufi and Khoshkhoh [2015] point out, the within-region change in sentiment should be driven by a *change* in party in the White House. If the first-stage equation really reflects a shift in sentiment, rather than a change in unobservables, I should not register a significant shift in house price expectations after the victory of an incumbent President in 2012. In Column 4 of Table 2.7 I therefore estimate Equation (9) again, to check whether there is a change in house price expectations around the 2012 presidential election. The coefficient on the interaction term is positive but not statistically significant. Also, its magnitude is three times smaller than the comparable estimate in Column 1 of Table 2.7, which instead presents the result of the same model for the 2008 presidential election.

The results of this section suggest that there was a shift in housing market sentiment occurring at the time when the Democrats won the White House in 2008. This shift was positively correlated with the ex-ante political views of a state's population: Democratic-leaning states became more optimistic than Republican-leaning states. This shift does not seem to be due to changes in housing or federal policies, or due to a generalized optimistic view about the economy. Rather, the results suggest that this is likely to be a pure shift in optimism affecting particularly the housing market. Therefore, conditional on a set of covariates, the interaction between partisanship and electoral timing appears to be a relevant and valid instrument to analyze the effects of a change in house price expectations on mortgage borrowing.

2.3.4 Second stage: expectations and mortgage leverage

Table 2.8 presents the analysis of the effects of housing sentiment on mortgage leverage ratios. Column 1 displays the simple OLS regression framework defined in Equation (8). Individual mortgages' loan-to-value ratios (LTVs) are regressed on a set of covariates, including different measures of sentiment in the same quarter/state cell.

House price expectations are positively correlated with mortgage borrowing (the estimated elasticity is +0.11, significant at the 5% level). Interest rate expectations are also positively and significantly correlated with LTVs, albeit the estimation is quite imprecise. This sign is coherent with what could be expected, since the Freddie Mac Single Family Loan Level Dataset includes only fixed-rate mortgages, and therefore an expected rise in interest rates should lead consumers to buy/refinance when rates are more favorable. On the other hand, expectations about inflation, personal income or opinions about government policy display no significant correlation with the dependent variable. This is an important result, because my measure of opinions about government policy is the same as the one used by Mian, Sufi and Khoshkhoh [2015] in their analysis. Perceptions of the government have no effect on household consumption decisions (Mian, Sufi and Khoshkhoh [2015]); similarly, I find no correlation between this measure and mortgage leverage. This should at least partially confirm the validity of the instrumental variable strategy proposed here. The 2008 presidential election might have shifted both government and housing sentiment however the former seems to have no direct implications on household consumption decisions. Therefore, it might be possible to assume that any significant effect estimated in the second stage relationship reflects a shift in housing, rather than a shift in government sentiment.

Column 2 runs a similar model to Column 1, only excluding other sentiment variables and including some aggregate controls for state-level time varying variables which might be correlated with both expectations growth and changes in mortgage leverage.⁴⁷ The coefficient associated with house price expectations is slightly smaller, but still positive and significant at the 10% level.

Finally, Column 3 shows the results of the IV estimation. I instrument the state-level change in house price expectations with the interaction between state-level political leanings and election timing (as described by equations 9 and 10). These

⁴⁷. These are: the state level GDP (logs); unemployment rate; average change in house prices over the past year; average property taxes (logs); population growth; number of building permits (logs); average wages in the construction sector (logs).

results are directly comparable to the OLS estimation presented in Column 2, since both include the same set of controls and the sample size is identical.⁴⁸ I can reject the null at the 1% level (with an estimated elasticity of +0.6). The loan-level characteristics have the expected signs: longer mortgages have higher LTVs and homeowners with higher credit scores are generally less heavily in debt. Interest rates are positively correlated with loan-to-value ratios, which is probably explained by the credit risk associated to lending a larger proportion of a property's appraisal value. First-time homebuyers, on average, receive less lending on similar properties (probably reflecting a shorter credit history). Loans on investment properties are also about 2.5 percentage points lower than loans on owner-occupied properties.

This result indicates that one standard deviation increase in state-level house price expectations (1.7 percentage points) generates an increase in individual-level loan-to-value ratios worth 1 percentage point, or 6% of a standard deviation in the dependent variable.

2.3.5 Different types of mortgages

The mortgages included in Freddie Mac's portfolio serve different purposes. About 37% are loans to buy a property and the remaining 63% of the loans are refinancing mortgages. This type of loans are meant to extract equity from *already occupied* properties. Among refinancing mortgages, a further distinction needs to be made between "cash-out" loans and "non-cash-out" loans. The former type is free from any specific purpose, however the latter is a mortgage that has the intent to pay off existing mortgage and house-related debt.⁴⁹ The three types of loans are roughly equally represented in Freddie Mac's portfolio between 2007 and 2014: purchase mortgages constitute 37% of the total, cash out mortgages 28%, and non-cash-out ones about 34%.

However, these three types of mortgages might be affected by house price expectations in different ways. Borrowing with the expectation of house price increases makes sense if households are liquidity constrained, desire higher consumption (whether housing-related or not), and bet on home appreciation to pay off a part of their debts in the near future. This type of logic is less likely to apply to

48. Column 2 also excludes observations recorded in the election quarter (q4 2008).

49. The loan is limited to being used to: pay off the first mortgage, regardless of its age; pay off any junior liens secured by the mortgaged property, that were used in their entirety to acquire the subject property; pay related closing costs, financing costs and prepaid items; disburse cash out to the borrower (or any other payee) not to exceed 2% of the new refinance mortgage loan or \$2,000, whichever is less.

households who are opening a second mortgage in order to pay off existing debts. In this latter case, the decision to refinance is most likely due to the desire to change the mortgage conditions (length, or structure of the interest rates) due to changes in policy or to unforeseen circumstances, such as the loss of employment.

Table 2.9 explores the effect of house price expectations on these three different types of mortgages. House price expectations display an elasticity of +0.62 on purchase mortgages (Column 1, significant at the 5% level), but the effect is almost double on cash-out refinancing mortgages (+1.21, Column 2, significant at 1% level). On the other hand, the coefficient associated with house price expectations on non-cash-out refinancing mortgages is substantially smaller (+0.34, Column 3) and statistically insignificant. Aggregate level controls support the idea that non-cash out refinancing mortgage borrowing is driven mainly by negative circumstances, rather than by speculation about the future of the housing market. One standard deviation increase in the unemployment rate is correlated with an increase in non-cash-out refinancing borrowing worth 0.46 percentage points (2.7% of a standard deviation, Column 3). Average wages in the construction sector (a proxy for the average level of wages in a state) also negatively correlate with loan-to-value ratios.

Finally, I address the concern of unobserved fundamental shock that could be correlated with state ideology and emerges *after* the election. As a robustness check, in Table 2.10 I estimate Equation (10) again, including only the second and third quarter of 2008, and the first two quarters of 2009 (2008q4, or the election quarter, is excluded). Analyzing the effect of expectations in a temporal span so close to the election helps to further reduce the concern that my results could be driven by unobservable shocks affecting more progressive states in the post-electoral period.

Column 1 shows that between 2008q2 and 2009q1, leverage on purchased properties was not significantly affected by changes in house price expectations. The coefficient (+0.08) is positive, but not statistically significant. The same is true of mortgage refinancing to pay off existing debt (Column 3). On the other hand, the estimated elasticity on cash-out refinancing mortgages (+1.76) is statistically significant at the 5% level (Column 2). The magnitude of this effect implies that an increase one standard deviation increase in short term house price expectations between 2008q2 and 2009q2 (excluding the quarter of the election) generated an increase in cash-out refinancing mortgage leverage ratios worth 2.9 percentage points, or 18% of a standard deviation in the dependent variable.

Overall, the effects of a shift in housing sentiment on mortgage leverage are significantly positive and robust to different specifications. If consumers on average expect house prices to rise in the near future, the individual leverage ratio increases, and as a consequence so does the leverage ratio of the aggregate economy.

2.4 Conclusions

In this chapter, I document the pattern of house price expectations formed by American consumers in the aftermath of the 2007 financial crisis. I show that expectations are heterogeneous across the population and that they contain a component of systematic extrapolative bias which is inconsistent with full-information rational expectations theory. Finally, motivated by the role that mortgage leverage had in the 2007-2008 financial crisis, I study whether house price expectations might be considered a fundamental driver of mortgage borrowing and lending behaviour. By exploiting an exogenous shift in housing sentiment that occurred after the 2008 presidential election, I show that a change in house price expectations has substantial effects on mortgage leverage, which increases whenever there is an expected increase in home equity.

These results have interesting implications. In the presence of cognitive biases like the representativeness heuristic, the asset markets may be subject to endogenous excess volatility (Gennaioli, Shleifer and Vishny [2015]; Bordalo Gennaioli and Shliefer [2016]). Since houses are also effectively used as collateral for loans (DeFusco [2015]; Mian and Sufi[2011]), the dynamic relationship between asset prices and mortgage leverage may exacerbate the cycle even further and pose substantial risks to financial and macroeconomic stability (Jordà et.al [2015a,b]; Jordà et.al [2016]). The extrapolative heuristic in house price expectations, combined with the role that such expectations have on mortgage leverage decisions, may therefore help explain the relationship between the housing cycle and the consumption cycle, and perhaps even the slow recovery after the Great Recession (De Nardi [2014]).

A theme that connects recent empirical research on expectations is the role of personal experiences, as opposed to public information, in shaping individuals' beliefs about the future. For example, cross-sectional "contagion" among peers affects the average consumer's decisions about investing in the housing market more than objective economic data (Bayer, Mangum and Roberts [2016]; Bailey et.al.[2016]).⁵⁰ The evidence on heterogeneity in expectations presented in this

chapter may indeed arise from private information, but also from differential access to public information or even from different forecasting models: my work can be improved by distinguishing between these alternative channels. Analysing the drivers of expectations' heterogeneity may constitute the empirical basis upon which to develop a new theory of how consumers are likely to react to news, a theory that does not rely on the assumption that all agents are perfectly informed and efficient in their forecasting methods. Indeed, this chapter shows that full-information rational expectations theory may not be entirely capturing the dynamics of the housing market.

Developing a more realistic model of how beliefs are formed seems particularly important because expectations are more than just noise: the evidence I present here suggests that they directly affect the credit cycle. To further refine my results it would be interesting to disentangle whether house price expectations affect primarily lending standards or mortgage applications. Nevertheless, the analysis of consumers' expectations appears to be a useful tool in the identification of emerging asset and credit bubbles.

50. Kuchler and Zafar [2015] show that personal experiences also strongly guide the individual's perception of the labor market. Such perceptions might affect not only job search efforts, but also labor demand and investment, if employers form expectations similarly to their employees, suggesting the possible existence of unemployment cycles (Eeckhout and Lindenlaub [2015]).

Figures

Figure 2.1. Expectations on income, inflation and house price growth rates at the 1-year horizon. US average, 2007-2014. Source: Michigan Survey of Consumers.



Figure 2.2. House prices forecast errors, percentage points. US average by quarter. Sources: Michigan Survey of Consumers; FHFA quarterly repeated sales index.

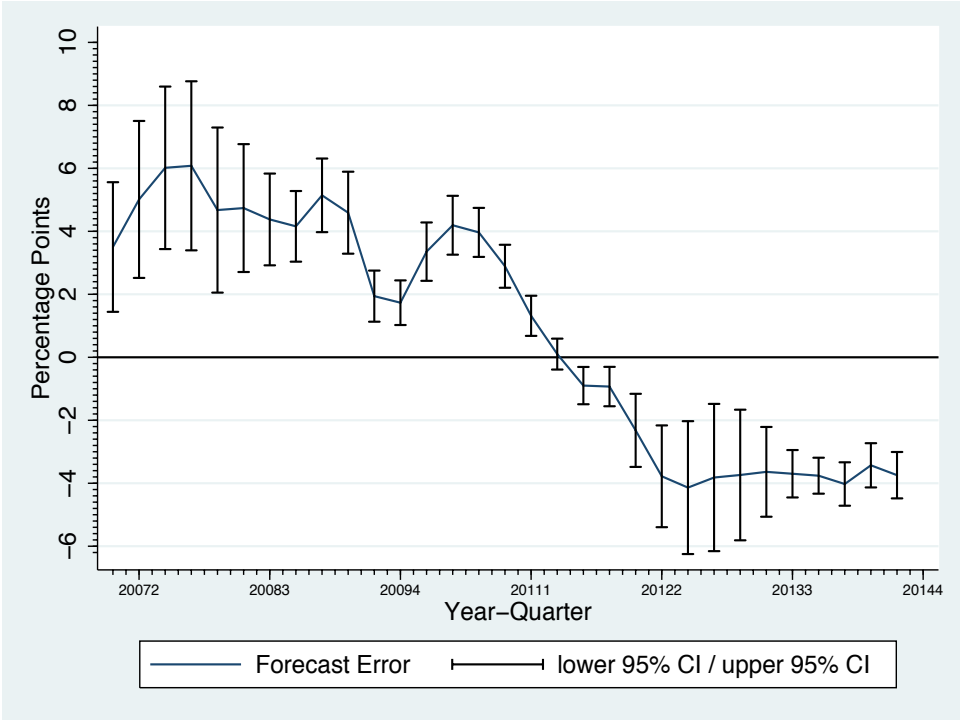
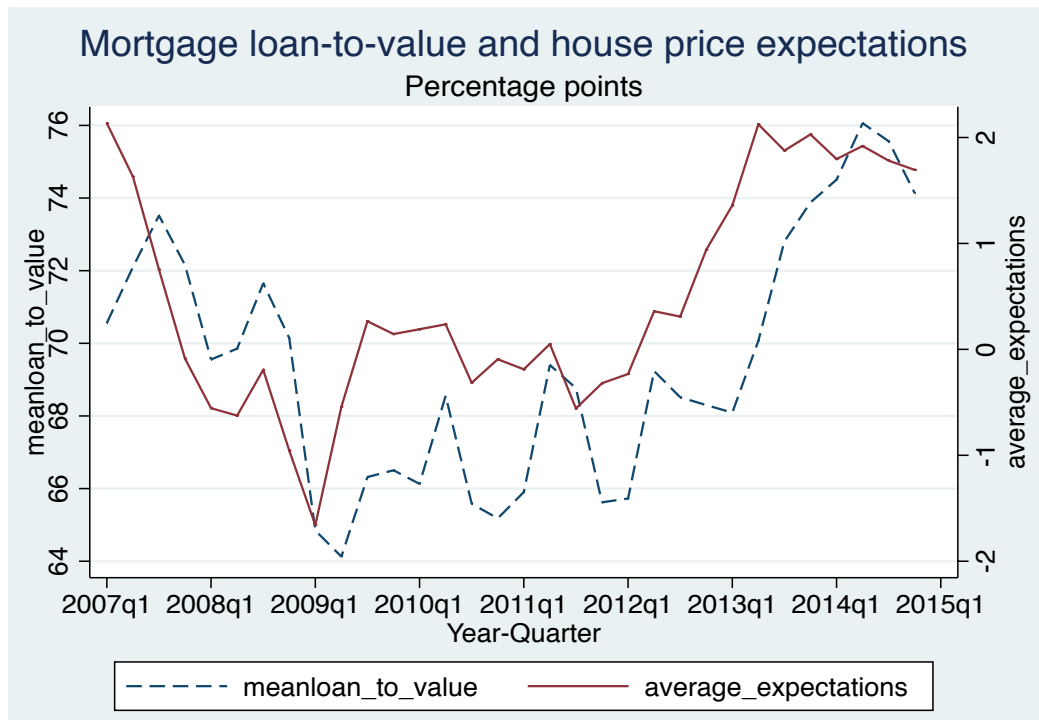


Figure 2.3. Correlation between average mortgage loan-to-value ratios and average expectations about 1-year ahead house price growth. Percentage points; US averages by quarter. Sources: Michigan Survey of Consumers; Freddie Mac Single Family Loan-Level Dataset.



Tables

Table 2.1 Descriptive statistics

Summary Statistics, 2007-2014						
Variable	Units	Obs	Mean	Std.Dv	Min	Max
<i>Panel A</i>		Michigan Survey of Consumers				
HH Income	Yearly, US \$	48896	84074.71	73331.61	2400.00	500000.00
Age	Years	52304	55.17	56.59	18.00	97.00
Male	Dummy	52548	0.46	16.59	0.00	1.00
Married	Dummy	52548	0.67	0.46	0.00	1.00
Adults	#	52548	1.90	0.70	1.00	5.00
Children	#	52521	0.61	1.03	0.00	5.00
College Educ.	Dummy	52548	0.52	0.50	0.00	1.00
Stock Owner	Dummy	52548	0.70	0.45	0.00	1.00
Exp. Hprice 1 Y	% growth	36404	0.34	5.30	-25.00	25.00
Exp. Income 1Y	% growth	50371	1.77	14.09	-50.00	95.00
Exp. Inflation 1Y	% growth	47047	3.84	4.15	-10.00	20.00
Forecast Error Hprice	%	35137	0.85	7.04	-30.85	25.00
Absolute Value For.Err.	%	36313	5.12	4.85	0.00	30.85
<i>Panel B</i>		Freddie Mac Single-Family Loan Level Dataset				
Loan to Value	%	399296	69.48	17.36	6.00	100.00
Length Mortgage	Years	399304	26.12	6.60	5.00	43.00
Credit Score	Points	399239	753.69	46.71	333.00	844.00
Interest Rate	%	399304	4.76	1.06	2.25	9.13
First Time Buyer	Dummy	399304	0.47	0.50	0.00	1.00
Investment Property	Dummy	399304	0.06	0.24	0.00	1.00
Purchase	Dummy	399304	0.37	0.48	0.00	1.00
<i>Panel C</i>		State-Level Controls				
Change House Price t-1	Percentage points	395324	0.00	0.02	-0.08	0.06
Unemployment Rate	Percentage points	395324	0.06	0.02	0.02	0.12
Population Growth	Percentage points	390182	0.01	0.01	-0.03	0.05
Building Permits	# per year	390182	33796.58	34736.81	536.00	176992.00
Wage Construction Sector	Monthly, US \$	389732	4312.46	637.01	2724.00	6756.00
Property Tax	Yearly, US \$	397938	2011.87	836.06	462.01	5346.01
Public Housing	Percentage/ pop.	397938	0.01	0.01	0.00	0.08
Rent Subsidies	Percentage/ pop.	397938	0.01	0.07	0.00	1.00

Table 2.2. Determinants of individual-level expectations

VARIABLES	(1) Income t+1	(2) Inflation t+1	(3) House Price t+1
HH Income (log)	-1.654*** (0.190)	-0.527*** (0.048)	0.211*** (0.072)
Age	-0.310*** (0.052)	0.049*** (0.010)	-0.037** (0.014)
Male	0.812*** (0.166)	-0.503*** (0.072)	0.200** (0.077)
College Degree	1.613*** (0.236)	-0.348*** (0.061)	0.287*** (0.070)
Married	-0.555** (0.267)	0.200*** (0.067)	0.078 (0.069)
#Children	0.232* (0.134)	-0.038 (0.030)	-0.100** (0.038)
Stock Owner	0.171 (0.241)	-0.294*** (0.065)	0.246*** (0.083)
Negative Income Shock	-2.466*** (0.175)	0.856*** (0.052)	-0.855*** (0.067)
Change HPrice Lag(State)	-3.464 (11.395)	0.497 (2.353)	52.841*** (3.863)
Income Change (State)	0.121 (1.824)	1.575** (0.587)	0.478 (0.799)
Gini Coeff.(State)	17.562* (9.400)	-7.288** (2.987)	4.540 (4.881)
Unemployment Rate(State)	-41.759*** (12.617)	1.093 (3.514)	7.344 (6.786)
Constant	28.046*** (5.439)	12.140*** (1.649)	-3.096 (2.935)
State FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Observations	33,976	32,020	33,892
R-squared	0.049	0.072	0.063

This table shows how individual-level expectations depend on individual-level characteristics and aggregate dynamics. The source is the University of Michigan Survey of Consumers , 2007 to 2014, only homeowners are included. The dependent variable in column 1 is the expected % personal income change in one year; in column 2 is the expected inflation change in one year; in column 3 is expected house price growth in 12 months. Errors are robust to heteroskedasticity and clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1

Table 2.3. Determinants of individual-level house price forecast errors

VARIABLES	(1) Inaccuracy Forecast	(2) Forecast Error House Price	(3) Delta Forecast Error
Delta Personal Income			0.404** (0.158)
Delta House Price lag			0.677*** (0.100)
Age	0.018* (0.011)	-0.041** (0.019)	0.035 (0.021)
Male	-0.265*** (0.049)	0.246*** (0.086)	0.156 (0.108)
College	-0.144** (0.059)	0.321*** (0.094)	0.131 (0.154)
Married	-0.081 (0.088)	0.069 (0.068)	-0.215 (0.154)
#Children	-0.044 (0.030)	-0.127*** (0.042)	-0.016 (0.066)
Stock Owner	-0.300*** (0.077)	0.199 (0.127)	-0.020 (0.065)
Negative Income shock	0.259*** (0.078)	-0.743*** (0.079)	0.101 (0.155)
Delta Unemployment			0.033 (0.123)
HH Income (logs)	-0.252***	0.289***	
State FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Constant	6.805*** (0.699)	0.733 (1.065)	3.998** (1.916)
Observations	32,844	32,844	11,809
R-squared	0.262	0.317	0.098

This table analyzes individual level forecast errors on house prices from the panel component of the Michigan Survey of Consumers, 2007 to 2014. The dependent variable in Column 1 is the absolute value of the house price forecast error: the further away from zero, the higher the inaccuracy of the forecast. In Column 2 the dependent variable is the forecast error, defined as in Equation (6): a higher value implies excessive optimism with respect to future house price realizations. In Column 3 the dependent variable is the change in forecast errors between the two interviews (quarter $t+2$ and quarter t). A higher value implies an increase in over-optimism. Standard errors are robust to heteroskedasticity and clustered at the state level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2.4. First-stage: house price expectations and political outcomes

VARIABLES	(1)	(2)	(3)
	Exp. H. Price t+1	Exp. H. Price t+1	Loan To Value Ratio
	OLS First Stage	OLS First Stage	OLS Reduced Form
DemShare(08)*PostElection(08)	0.103*** (0.028)	0.071*** (0.014)	0.051** (0.020)
Dem.Share(08)	-0.057** (0.024)	0.001 (0.033)	-0.322*** (0.041)
Post Election(08)	-5.977*** (1.265)	-4.117*** (0.774)	5.791*** (1.001)
Aggregate Income		1.201 (1.250)	-1.496 (1.380)
Unemployment_rate		-2.501 (5.883)	11.342 (8.270)
Delta HPrice Lagged		49.485*** (2.983)	10.541*** (3.693)
Income Borrower	-0.013** (0.005)	-0.023*** (0.005)	3.261*** (0.244)
Loan Length	-0.003 (0.002)	-0.002 (0.002)	0.394*** (0.014)
Credit Score	-0.000* (0.000)	-0.000 (0.000)	-0.045*** (0.002)
Interest Rate	0.082 (0.066)	0.046 (0.047)	3.081*** (0.210)
First Time Buyer	-0.016 (0.020)	-0.012 (0.023)	-1.501*** (0.179)
Investment Property	-0.041** (0.020)	-0.035* (0.019)	-1.714*** (0.173)
Purchase (VS Refinancing)	0.052** (0.020)	0.109*** (0.017)	11.287*** (0.577)
State FE	yes	yes	yes
Year FE	yes	yes	yes
Constant	3.661*** (0.969)	-30.713 (34.354)	98.129** (38.492)
Observations	380,389	380,389	380,612
R-squared	0.399	0.492	0.269

This table shows first-stage and reduced form relationships between electoral outcomes, house price expectations and loan-to-value ratios. The dependent variable in Columns(1)-(2) is the quarter/state average of one-year house price expectations. In Column.(3) the dependent variable is the individual-level mortgage loan-to-value ratio. All models include both state and year fixed-effects; standard errors allow for heteroskedasticity and are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.5: First-stage robustness: post-electoral policy change?

VARIABLES	(1) Public Housing	(2) Rent Subsidies	(3) Property Taxes	(4) Building Permits	(5) Wages Construction Sector
	T+N	T+N	T+N	T+N	T+N
Dem Share (08)*PostElection (08)	-0.000	0.002	0.003	0.009	0.000
	(0.000)	(0.002)	(6.576)	(114.537)	(3.005)
Quarter FE	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes
Observations	1,410	1,410	1,410	1,370	1,369
R-squared	0.012	0.035	0.361	0.368	0.805
Number of fips	48	48	48	47	47

This table shows the relationship between electoral outcomes and state-level housing policy at t+N. The dependent variables are averages at the state level, and they originate from the March CPS, the Bureau of Labour Statistics and the Census Bureau, which are: Percentage of citizens living in public housing (Column 1); Percentage of renters who receive rent subsidies (Column 2); Average property taxes on residential housing (Column 3); Yearly building permits issued by the local authorities (Col.4); Average Wages in the construction Sector (Column 5). Time fixed-effects are also included. Errors are robust to heteroskedasticity and are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1

Table 2.6: Robustness: other expectations?

VARIABLES	(1) Exp. Personal Income	(2) Exp. Inflation	(3) Exp. Interest Rates	(4) Positive View of Government	(5) Exp. House Price
Dem Share(08)*Post Election(08)	0.022 (0.027)	-0.039*** (0.008)	-0.003 (0.002)	0.008*** (0.001)	0.044*** (0.016)
Positive View of Government					1.453** (0.767)
Expect. Int Rate 1					-0.256 (0.301)
Expect. Income 1					0.045*** (0.014)
Expect Inflation 1					-0.055 (0.046)
Aggregate Controls	yes	yes	yes	yes	yes
Quarter Fixed effects	yes	yes	yes	yes	yes
Observations	1,458	1,457	1,458	1,458	1,451
R-squared	0.060	0.288	0.334	0.294	0.254
Number of fips	49	49	49	49	47

This table looks at the relationship between electoral outcomes and sentiment measures different from house prices. The dependent variables are compiled from the Michigan Survey of Consumer Sentiment (2007-2014) as state/quarter averages: Expectations about personal income percentage growth in one year (1); about inflation rates growth in one year (2); about whether interest rates will go up/down in one year (3); about whether the respondent has a positive view of the government's policy(4). Col.(5) regresses house price expectations on all other sentiment measures and the interaction term between post electoral quarters and democratic vote shares in 2008q4. All columns include quarter fixed effects. Aggregate controls include the sum of log income; unemployment rates; average house price growth in the previous three quarters; average wages in the construction sector; number of building permits; population growth (year on year). Errors are robust to heteroskedasticity and are clustered at the state level.*** p<0.01, ** p<0.05, * p<0.1

Table 2.7: Robustness: placebo tests

VARIABLES	(1) ExpHPrice1	(2) ExpHPrice1	(3) ExpHPrice1	(4) ExpHPrice1
DemShare(08)*PostElection(08)	0.072*** (0.015)			
DemShare(08)*2010 onwards		0.028* (0.016)		
DemShare(08)*2011 onwards			-0.044 (0.029)	
DemShare(12)*PostElection(12)				0.022 (0.014)
State FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Constant	2.322*** (0.256)	-2.056*** (0.391)	0.117 (0.245)	2.316*** (0.250)
Observations	1,451	1,076	888	1,451
R-squared	0.237	0.230	0.188	0.231
Number of fips	47	47	47	47

The dependent variable is the state/quarter averages of house price expectations recorded by the Michigan Survey. Column 1 shows the response of house price expectations to the instrument (the democratic voting share in the 2008 election interacted with a dummy indicating the post-electoral period). Columns 2 and 3 display whether the voting share for the Democratic party in the 2008 elections shifted house price expectations also in subsequent years. Column 4 shows the response of house price expectations to an alternative instrumental variable (the interaction between the voting share for the Democratic Party in the 2012 election with a dummy indicating the post-2012 electoral period). Errors are clustered at the state level.*** p<0.01, ** p<0.05, * p<0.1

Table 2.8. Second-stage: aggregate sentiment and individual-level loan-to-value ratios

VARIABLES	(1) LTV OLS	(2) LTV OLS	(3) LTV 2SLS
Exp HPrice1	0.119** (0.056)	0.079* (0.047)	0.604*** (0.231)
Exp inflation 1	-0.007 (0.027)		
View Govt.	0.061 (0.488)		
Exp Income 1	-0.008 (0.011)		
Exp.Int Rate 1	0.581** (0.232)		
Length Mtg.	0.406*** (0.014)	0.317*** (0.013)	0.317*** (0.013)
Credit Score	-0.038*** (0.002)	-0.043*** (0.002)	-0.043*** (0.002)
Interest Rate	3.723*** (0.278)	4.972*** (0.258)	4.971*** (0.249)
First Time Buyer	-1.336*** (0.180)	-1.416*** (0.181)	-1.399*** (0.173)
Investment Property	-3.309*** (0.131)	-2.433*** (0.181)	-2.432*** (0.180)
State FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Aggregate Controls	No	Yes	Yes
Constant	63.360*** (1.824)	147.098*** (32.547)	154.051*** (32.740)
First stage F stat.			20.21
Observations	394,958	372,755	372,755
R-squared	0.257	0.273	0.272

This table displays the effects of a change in house price expectations on individual-level mortgage loan-to-value ratios (measured by the micro data included in Freddie Mac's single Family Loan-Level dataset). Column 1 shows the results of an OLS estimation which includes controls for other expectations (income, inflation, interest rates and perception of the government). Column 2 displays an OLS estimation which excludes q3 2008 and expectations other than house price forecasts, but includes state-level controls such as: state-level GDP; unemployment rates; change in house prices in the previous four quarters; population growth; property taxes (average); average wages in the construction sector; number of issued building permits. Column 3 shows the results of the 2SLS estimation. Errors are clustered at the state level. ***p<0.001, p<0.05, *p<0.1.

Table 2.9: Different mortgages

VARIABLES	(1)	(2)	(3)
	Loan To Value 2SLS Purchase	Loan To Value 2SLS Refinancing Cash Out	Loan To Value 2SLS Refinancing Not Cash Out
Exp. HPrice 1	0.632** (0.311)	1.213*** (0.325)	0.343 (0.432)
Aggr. Income (log)	-2.395 (2.003)	-2.060 (1.821)	-1.114 (2.122)
Unemployment Rate	10.724 (10.373)	-6.490 (13.663)	23.445*** (9.007)
Change HPrice Lag	-55.538*** (17.806)	-35.804* (20.591)	16.592 (24.406)
Pop. Growth	31.898** (16.071)	36.239*** (13.750)	-15.938 (14.136)
Property taxes (log)	-0.147 (0.674)	0.024 (0.802)	1.661** (0.653)
Wages Construction (log)	-1.196 (2.646)	-9.124*** (2.180)	-9.118*** (2.719)
Building Permits (log)	0.122 (0.471)	-1.112* (0.641)	0.442 (0.785)
Dem Share (08)	0.005 (0.047)	-0.045 (0.068)	-0.128* (0.071)
Post Election (08)	19.651*** (2.506)	15.200*** (2.601)	16.353*** (2.282)
State FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
HH Controls	Yes	Yes	Yes
Constant	127.606** (60.042)	200.218*** (50.841)	172.070*** (64.591)
Observations	138,789	105,009	131,199
R-squared	0.129	0.112	0.155

This table analyses the relationship between house price expectations and loan-to-value ratios for different types of mortgages. Column 1 limits the analysis to mortgages which have the purpose of acquiring a property; Column 2 instead looks at cash-out refinancing mortgages only, where the borrower draws equity out of a property they already possess, and the loan has no specific purpose; Column 3 looks at refinancing non-cash-out mortgages. The model is a 2SLS estimation, identical to the one in Table 2.8, and includes the same set of controls. Standard errors are robust to heteroskedasticity and clustered at the state level.*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2.10. Robustness: sentiment change around the election

VARIABLES	(1) Loan To Value 2SLS Purchase	(2) Loan To Value 2SLS Refinancing Cash Out	(3) Loan To Value 2SLS Refinancing Not Cash Out
Exp. HPrice 1	0.084 (0.898)	1.768** (0.725)	0.067 (0.475)
Dem Share (08)	0.015 (0.011)	-0.214*** (0.029)	-0.079*** (0.022)
Post Election (08)	6.668*** (0.655)	2.967*** (0.853)	1.792** (0.717)
State FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
HH Controls	Yes	Yes	Yes
Constant	58.372*** (3.168)	64.981*** (4.026)	96.137*** (3.955)
Observations	15,214	15,234	18,912
R-squared	0.105	0.092	0.133

This table analyses the relationship between house price expectations and loan-to-value ratios for different types of mortgages between quarter 2008q2 and 2009q2, excluding the quarter when the election took place. Column 1 limits the analysis to mortgages which have the purpose of acquiring a property; Column 2 is limited to cash-out refinancing mortgages only; Column 3 looks to refinancing non-cash-out mortgages. The models include all controls included in Tables 2.8 and 2.9. Standard errors are robust to heteroskedasticity and clustered at the state level.*** p<0.01, ** p<0.05, * p<0.1.

Appendix A2.1: Control variables

Aggregate, state-level variables control for the characteristics of the housing market and the general economy in a given state and quarter. In both sets of estimations (on expectations and on mortgage leverage) I include past house price growth, measured at the state level, defined as the growth rate in the previous four quarters (Federal Housing Finance Agency repeated sales index). Some models control for time-varying fundamental shock (income growth, unemployment rates, homeownership rates (from the March CPS) and for changes in local housing policy (average property taxes, percentage of residents living in public housing, percentage of residents paying lower rent due to government subsidies). These variables are derived from the March CPS.

Both expectations and mortgage markets are likely to be affected by changes in regulation or other factors restricting housing supply, such as higher building costs. I proxy for production costs using average wages in the construction sector (Bureau of Labour Statistics, NAICS 23). The changes in the restrictiveness of regulation are proxied by the yearly number of building permits issued in a given state, which are here used as a measure of housing supply elasticity, as in Kahn [2011].⁵¹

Population growth and migration can also have important implications for the housing market equilibrium. I therefore also control for (yearly) changes in the number of residents using data from the US Census. Data on both 2008 and 2012 electoral results by state comes from the Federal Election Commission dataset. Panel C of Table 2.1 provides the summary statistics relative to these state-level variables.

⁵¹. Saiz [2010] provides a more precise measure of housing supply elasticity. This is however only available for certain metropolitan areas, rather than states.

CHAPTER 3

Fiscal Adjustment and Social Unrest in IMF-Supported Programs

3.1 Introduction

This chapter shifts the focus to public debt and deficits. The 2007-2008 financial crisis affected the public balance sheets of many developed economies. This was due in part to the direct cost of bank rescues, and in part to the cost of the fiscal stimuli which many governments initially implemented in order to sustain aggregate demand in face of the credit crunch and the recession. These events, coupled with pre-existing fiscal imbalances, triggered the onset of a public debt crisis in the periphery of the Eurozone in 2009.

The Eurozone crisis sparked a widespread debate on the determinants and the consequences of fiscal consolidation. Technical arguments regarding the magnitude and direction of fiscal multipliers dominated a part of the academic literature (Alesina and Ardagna [2010], Eggerston and Krugman [2012]) while another side of the debate was concerned with the potential social costs of public budget cuts (Woo et al.[2013]; Kentikelenis et.al [2011, 2014]). However, a common denominator of this literature is to describe fiscal policy design in times of crises as a purely technical matter, which is largely independent from the social and political circumstances of a given country. This chapter studies instead the design of fiscal consolidation programs and provides evidence that fiscal policy choices in times of crises might be strongly driven by political, rather than economic, considerations.

Public debt crises often trigger the intervention of international financial institutions, such as the International Monetary Fund (IMF). IMF agreements are a typical example of situations in which domestic political constraints are assumed to cease to matter, because the Fund is not accountable to a country's electorate.

In this chapter, I present a new panel dataset on fiscal conditionality attached to all International Monetary Fund (IMF) lending agreements taking place between 2002 and 2012. This dataset provides a novel source information on *planned* fiscal

policy efforts, spanning 91 countries and ten years.⁵² I use this dataset to analyse whether the domestic politics of borrowing countries affect fiscal reforms included as conditions in IMF agreements. My results indicate that IMF conditionality is susceptible to domestic political constraints of recipient countries. In particular, fiscal conditionality is subject to a political force often neglected by public choice theory: extra-parliamentary opposition, or the threat of social unrest.

Models in political economics traditionally explain fiscal discipline, or its absence, as a consequence of a country's formal political institutions, such as electoral or fiscal rules.⁵³ However, recent literature suggests that governments might be delaying or avoiding public budget consolidation because they might fear the welfare loss that ensues from demonstrations, riots, and strikes that typically follow public budget cuts (Ponticelli and Voth [2011]; Tabellini and Passarelli [2013]). IMF arrangements, however, need not be subject to the same constraints, as the IMF enjoys a large degree of bargaining power and could, in principle, impose policies on borrowing countries in exchange for lending.

To explain how domestic conflict restricts policy choices in the context of IMF arrangements, I situate my results within the bargaining framework proposed by Putnam [1988], known as the theory of two-level games. This theory predicts that international negotiations can be shaped by the relative preferences of each negotiator's domestic constituency. Facing binding domestic constraints can constitute a bargaining advantage if *ex-post* defection from the agreement is costly for both parties. In this framework countries with stronger domestic opposition may use this constraint as a bargaining tool to obtain more lenient fiscal conditionality within IMF programs, all else held equal.

Building on this theoretical framework, this chapter's main contribution is to show that the threat of social unrest is indeed a political force capable of restricting fiscal policy choices in the context of IMF arrangements. My results indicate that two additional episodes of civil unrest, in each of the three years prior to the negotiations, reduce the planned fiscal consolidation efforts by 1.2 percentage points of GDP, or 21% of a standard deviation in the dependent

52. To my knowledge, the only other data source collecting a cross-country panel of fiscal reforms is Devries et al. [2011], who focus only on OECD economies. However, their dataset is unsuitable to study the political determinants of fiscal policy *decisions*. This is because Devries et al. [2011] measure fiscal adjustment as the actual, *ex-post*, budget balance change. Hence, there is no obvious way to discriminate between a government's intentions and the program's implementation. Instead, I collect data on *planned* fiscal consolidation, rather than on the actual fiscal consolidation *outcomes*.

53. The seminal model by Alesina and Drazen [1991] focuses on the war of attrition between conflicting interest groups in society. For a thorough review of the literature on the topic see Alesina and Perotti [1995] and Alesina and Passalacqua [2015].

variable. This is robust to controls for the economic circumstances of the country; to political variables which previous literature has shown to be relevant in determining IMF conditionality; and to some potential selection issues.

I also study how social unrest relates to qualitative targets in the realms of public wages, welfare expenditure, or labor market reforms. These reforms are also very likely to trigger public opposition, as suggested by Tabellini and Passarelli [2013]. Countries facing a stronger history of social unrest also receive fewer qualitative conditions in these policy realms. Two additional episodes of civil unrest, in each of the three years prior to the negotiations, reduce the number of these conditions by 0.25 (a fourth of a condition), or 11% of a standard deviation in the dependent variable. The threat of social unrest, however, has no negative effect on reforms that are less likely to trigger public protests, such as trade and financial sector reforms.

This chapter draws inspiration from two main strands of literature. The first relates to the political economy of IMF conditionality, particularly recent work focusing on how the domestic institutional settings of borrowing countries can explain the variation in IMF's programs and conditions.⁵⁴ Stone [2008] is one of the first proponents of this view, showing that the scope of IMF conditionality is less broad in democracies, and particularly under coalition governments. Stone [2008] interprets this result as evidence that under these political setups governments face more domestic opposition to reform and use this as a bargaining tool against the IMF to obtain less intrusive policy prescriptions. Stone's [2008] results are in line with the work of Caraway, Rickard and Anner [2012], who find that countries with stronger labour movements (measured by the relevance of trade unions) receive fewer conditions related to labour markets. They also find that this effect is particularly strong in democracies. Woo [2013] shows that conditions related to the public sector are influenced by domestic political constraints much more than financial sector conditions, which instead respond to supranational interests.

I contribute to the literature on the domestic political determinants of IMF conditionality by focusing on a particular a set of conditions (quantitative fiscal targets) that, to my knowledge, has not been previously studied. My results challenge a view, dominant in earlier literature, which predicts that if a government

54. This is defined by Stone [2008] as the "structuralist" view of IMF conditionality. This approach is complemented by the "realist" view, which focuses more on the geopolitical pressures conditioning the IMF (Copelovitch [2010], Gould [2003; 2006]) and by the "bureaucratic" view, which focuses on the internal power struggle among IMF staff (Dreher [2004]).

is more constrained domestically, it should use the IMF to enforce costly reforms on its citizens, blaming the responsibility on the Fund's officials (Vaubel [1986], Przeworski and Vreeland [2000], Vreeland [2003]). Instead, I show that the potential for extra-parliamentary opposition reduces the extent of conditionality within IMF programs, consistently with the results Caraway Rickard and Anner [2012] present regarding the negative relationship between labour market conditionality and relevance of trade unions.

My results also fit in the broad literature relating the political economy of fiscal policy. Recent contributions, in particular, point to the relationship between fiscal adjustment and extra-parliamentary opposition. Ponticelli and Voth [2011] show how fiscal austerity programs are very likely to cause a large extra-parliamentary opposition in the form of riots, protests, and strikes. Tabellini and Pasarelli [2013] claim that the fear of social unrest, rather than electoral outcomes, might be what prevents governments from engaging in fiscal consolidation. This is particularly plausible as governments seem not to be punished by voters after implementing austerity programs (Alesina, Carloni Lecce [2012]). I contribute to this literature by providing empirical evidence that the threat of social unrest can indeed be a powerful force in explaining fiscal consolidation, or its absence. Furthermore, I show that the electoral cycle enhances this result; democracies are more responsive than autocracies to the threat of unrest, and governments facing imminent elections even more so.

This chapter proceeds as follows. Section 3.2 provides an overview of Putnam's theory of two-level games, which constitutes the theoretical basis to understand how social unrest can influence fiscal policy in the context of IMF agreements. Section 3.3 presents the dataset and the empirical methodology, Section 3.4 presents results on fiscal conditionality, and Section 3.5 presents results on structural reforms. Section 3.6 concludes.

3.2 Theoretical framework: conditionality as a two-level game

Putnam's [1988] theory of international negotiations features two "chief negotiators" engaging in a bargaining process set in two stages, or "levels". Level I is international: each chief negotiator tries to reach an agreement with their foreign counterpart. Level II is domestic: each level I agreement has to be ratified from the respective domestic constituencies.

In order to model the interaction between the two levels, Putnam uses the concept of “win sets”. A negotiators’ win set is the set of all feasible level I agreements that would win the approval of domestic constituencies at level II. Hence, the *absolute* size of both win sets defines the first of Putnam’s hypotheses: larger win sets should make agreements more likely, as the win sets overlap more easily. This hypothesis cannot be tested in my data, however, since I observe the population of *agreed* IMF conditions. On the other hand, the outcome of the negotiating process depends crucially on the *relative* size of win sets. In particular, Putnam’s theory states that a relatively smaller win set can constitute a bargaining advantage, and tilt the balance of the negotiations in favor of the party facing the highest degree of domestic constraints.

When applying this framework to the bargaining process between countries and the IMF over conditionality it is necessary to assume that the IMF and the country’s representatives have different preferences over economic policy. To simplify, I assume that the government will try to minimize conditionality and maximize lending, and that the Fund will have the opposite strategy. The Fund might prefer more conditionality for a given amount of lending, as it views policy prescriptions as a form of collateral: conditions are meant to safeguard IMF resources by ensuring that the country’s balance of payments will be strong enough to permit it to repay the loan.⁵⁵ The Fund might also be maximizing conditionality in order to expand its influence on the international stage, and therefore its staff’s power and prestige (Vaubel [1986], Vaubel Dreher and Soylu [2007]). Major stakeholders and other lenders involved in the negotiation process might also affect the Fund’s preferences over conditionality: this can sometimes soften the Fund’s stance in exchange for political positions that are favorable to powerful countries (Stone [2008]; Dreher and Jensen [2007]; Dreher Sturm and Vreeland [2009a,b]). However, external stakeholders might also desire more conditionality in exchange for lending (Gould [2003, 2006]). Holding these influences constant, it seems reasonable to assume that the Fund has a taste for more stringent conditionality, mainly because it could assure the swift payback of the loan tranches.

On the other hand, a government might try to minimize conditionality for a given amount of lending. This is because agreeing to fewer conditions may imply higher flexibility in the path to economic recovery and the government might value this

55. IMF [2016] factsheet on conditionality.

flexibility (Kahler [1993]). Also, the government might be concerned about the welfare loss associated with riots and protests that may result from proposing ambitious reforms to the citizens.⁵⁶

Given a budget constraint upon which both subjects agree, and assuming that the IMF's and the governments' priorities diverge, the set of feasible agreements can be represented by the segment $\overline{F^*G^*}$ in Figure 3.1. The segment $\overline{F^*G^*}$ depicts all the possible agreements satisfying both win sets and the budget constraint (which is defined as the combination of lending and conditionality which is capable of solving the country's macroeconomic imbalance). F^* is the preferred outcome for the Fund, as it minimizes lending and maximizes conditionality; G^* is the preferred outcome for the government, for the opposite reason.

Ilida [1993] formalizes the relationship between the size of the respective win sets and the distribution of the negotiation outcomes. In Ilida's model there are infinite periods. In each period the players, G and F , take turns proposing an agreement over the contract curve $\overline{F^*G^*}$ (Figure 3.1), which is a proposal over a given amount of lending (l) and a set of policy prescriptions (c). Assuming that citizens have to decide whether or not protest against the proposed reforms (lc), and that each citizen has a status quo payoff (s_i). Each citizen is indifferent between the status quo and the payoff from the negotiations if and only if $s_i=lc$. If $lc \geq s_i$ citizen i decides not to protest. The key to the agreement is then going to be the median citizens' evaluation of the status quo (s_m).⁵⁷ To simplify the analysis, this model assumes that both the government and the Fund have perfect knowledge of the level of domestic constraints in the borrowing country. Also, it assumes that the governments' interests align with that of its constituency, and that the constituency is homogeneous in its demands (there are no conflicting interest groups lobbying the government about conditionality).⁵⁸

56. This assumption is in contrast with the setup proposed by Vaubel [1986] or Vreeland [2003], where governments use the IMF as a scapegoat, to weaken domestic opposition and enforce reforms on the population. If this were to be the case, we should observe more conditionality the higher the degree of domestic constraints; and more so in democracies than in autocracies. However this chapter as well as the work of Stone [2008], Caraway Rickard Anner [2012] and Woo [2013] contradict this earlier hypothesis.

57. Of course, the decision to participate in riots may be endogenous to other people's decisions. A more realistic model of the costly decision to participate in a riot as a consequence of economic policy is proposed by Tabellini and Passarelli [2013].

58. Regarding the common knowledge assumption, Ilida [1993] models the possibility of international asymmetric information, with one country having full knowledge of its constraints but not the other party. He finds that in most situations the bargaining power of the domestically constrained negotiator is still enhanced for the most part. On the other hand, Mo [1994] expands the analysis to a situation where the negotiators' interest are in conflict with at least some of the interests in her constituency and finds that domestic constraints can still be a bargaining advantage for the negotiator even in this setting.

Assuming common knowledge, then both negotiators know the real value of s_m , and they know that no agreement in which $lc < s_m$ can ever be ratified in the country. This is because in order for agreement $(lc, 1-lc)$ to be ratified, the median citizen needs to support it. Crucially, both the government and the Fund have positive discount rates, capturing the idea that not finding an agreement is costly for both parties. The country's motives are clear, as a delay in reaching an agreement might lead to greater costs associated with its macroeconomic imbalances. The Fund has a different set of incentives in finding an agreement as soon as possible. The internal dynamics of appointment strongly incentivize IMF negotiators to return from country missions with an agreed set of reforms (Kahler [1993]). Moreover, the delay might put pressure on the internal coalition brought together to support a given lending program, or negotiating position. To avoid unsettling internal dynamics, the IMF management might be willing to concede to some of the country's demands in order to reach an agreement as soon as possible (Kahler [1993]). Also, the threat of economic collapse may lead rich countries, who are the Fund's major stakeholder, to the rescue of their geopolitical allies. Finally, the Fund might be worried about the potential regional spillovers of a country's macroeconomic crisis. So it is reasonable to assume that both the government and the Fund will have an incentive to propose a solution which satisfies the other party as soon as possible.

Given this setup, the citizen's decision to riot affects the outcome of the negotiations through the welfare loss potentially associated with the protests. Anecdotal evidence suggests that citizen protests have strong-armed governments into revoking reforms in Italy, Portugal, and Greece, during the European debt crisis (Tabellini and Passarelli [2013]). Developing countries are not immune from this mechanism, as the case study of Jamaica displays (Kahler [1993]). If the citizens riot against a proposed reform, the government can be forced to repeal the agreement. This is costly for both the government and the IMF.

With such a mechanism at play, a higher degree of social unrest can change the equilibrium outcome of the bargaining game in favour of less fiscal adjustment, given the same amount of lending. In other words, as long as the Fund's discount rates are positive, a higher degree of domestic constraints constitutes a bargaining advantage for the government. Graphically, this can be represented by a change in the equilibrium outcome on the contract curve $\overline{F^*G^*}$ in Figure 3.1. If the government becomes more constrained domestically, its win set will shift from $[s_{m1}, G^*]$ to $[s_{m2}, G^*]$.

This result, formalized by Ilida [1993], relies on the assumption of a constant win-set on the side of the Fund. This is equivalent to assuming that the IMF always faces the same set of domestic constraints when dealing with a specific country. In reality, the Fund’s win set towards a specific borrower might shift, especially in response to changing geopolitical pressures. However, a shrinking win set on the side of the Fund will shift the outcome in its favor, but does not fundamentally change the intuition of the model.

The logic of two-level games provides a theoretical framework to investigate how domestic constituencies can influence international negotiations with the Fund, as Putnam originally suggested:

*“[...] if labor unions in a debtor country withhold necessary cooperation from an austerity program that the government has negotiated with the IMF, level II ratification of the agreement may be said to have failed; ex ante expectations about that prospect will surely influence the level I negotiations between the government and the IMF”.*⁵⁹

3.3 Data and methodology

This section describes the construction of the dataset and the definition of the two main dependent variables. It also discusses the construction of the main measure for social unrest and the methodological framework.

3.3.1 Data on IMF conditionality

The source of conditionality data is the Monitoring of Fund Arrangements (MONA) database, which was disclosed by the IMF to the public in 2009, permitting retroactive access to a large part of the lending agreements. The IMF disburses its loans to countries in subsequent tranches, over the course of an agreement that might span over several years. Most IMF programs are classified as conditional lending, in which the country is supposed to fulfil a set of policy prescriptions (conditions) before receiving the subsequent tranche of the loan. IMF representatives visit the borrowing country multiple times over the course of a program in order to discuss compliance with conditions that

⁵⁹. Putnam [1988], p.436

have been agreed in the past and to define new policies for the future. While the IMF might enjoy a large degree of bargaining power during these negotiations, the conditions are nevertheless an object of a bargaining process between the Fund's and the country's representative. The MONA database collects data on the outcome of this bargaining process, represented by the set of conditions agreed in each program review.

I focus only on programs taking place between 2002 and 2012. The MONA database provides information about the vast majority of agreements taking place in this time frame, while older agreements are available in the MONA archive (from 1993 to 2003). Due to their structure, the two datasets are not easily comparable. Choosing this particular time frame might be problematic because in the year 2000 the IMF launched an initiative to reduce the scope and extent of conditionality. Despite this, the number of structural conditions included in the average program had not changed substantially up to 2004 (Independent Evaluation Office [2008]). This suggests that the streamlining initiative was mainly interpreted by the Fund's staff as an invitation to change the composition of conditionality and to focus more on the IMF's area of core expertise (Independent Evaluation Office [2008]). This chapter focuses on rather homogeneous groups of conditions. Therefore, the choice of this specific time frame should not substantially condition the results.

Overall, I construct a dataset composed of 91 countries, 171 unique agreements, and 486 program/years. Figure 3.2 shows the geographical distribution of IMF agreements during this time frame. On average, countries in this dataset are part of an IMF agreement for 3.3 consecutive years, with the maximum number recorded for Tanzania, which has been under four consecutive IMF programs between 2002 and 2012. Table 3.1 contains the descriptive statistics for this sample. I construct a panel dataset, where each country is observed for a number of consecutive years. Since I only observe conditionality when a country is undergoing an IMF agreement, this is an unbalanced dataset from the point of view of the dependent variables (not each country/year cell is also a program/year cell).

3.3.2 Construction of conditionality variables

IMF conditions can be categorized along two dimensions. The first is their nature: conditions can be quantitative or qualitative. Quantitative conditions include Quantitative Performance Criteria (QPCs) and Indicative Targets (ITs). For

example, these can be quantitative targets on government budget balance or targets on the required level of international reserves. Qualitative conditions include Structural Performance Criteria (SPCs) and Structural Benchmarks (SBs), such as measures to improve the public or financial sector operations. Prior Actions (PAs) can belong to either category.

Conditions also differ in their degree of compulsoriness, as some are considered mandatory in order to pass a review, and as such they require a waiver from the Executive Board in case they are not met. Compulsory conditions include Prior Actions (PAs), which need to be met in order for the first tranche of the loan to be disbursed. Other mandatory conditions are Quantitative and Structural Performance Criteria (QPCs and SPCs). Indicative Targets (ITs) and Structural Benchmarks, while important, are intended as general markers to assess program performance, and compliance with them is not considered necessary in order to pass a review.⁶⁰

I compute two measures of conditionality from the MONA database. The first is the degree of fiscal consolidation required in a given country at the end of any program year.⁶¹ The MONA database provides data on the level of fiscal deficit required of country c at the end of year t expressed in billions (or millions) of National Currency. These are all defined as Quantitative Performance Criteria, and are therefore compulsory in order to pass a review. I collect this variable for all country/years for which it is available, and scale its nominal value the consolidated value of GDP recorded for country c in year t .⁶² The first of my dependent variables is therefore constructed as follows:

$$Consolidation_{ct} = Target_{ct} - Fiscal\ Deficit_{ct-1} \quad (11)$$

where $Target_{ct}$ is the expected deficit-to-GDP ratio included as an IMF condition at the end of year t and $Fiscal\ Deficit_{ct-1}$ is the deficit-to-GDP ratio at the end of year $t-1$.⁶³

60. This information is publicly available on the IMF website, under “conditionality factsheets”.

61. For a given country/year, there are multiple deficit targets, as the IMF runs multiple reviews per program/year. I decide to consider only the end-year target (the last one reported in the MONA database for each given country/year). This is for consistency with the macroeconomic explanatory variables, which are normally consolidated to reflect end-year values.

62. Some country/years are not program/years, and some program/years do not include a fiscal deficit condition. These are treated as missing values.

63. Fiscal deficit at $t-1$ is defined as the difference between the consolidated figures of government revenues and government expenditures at the end of year $t-1$, as included in the World Economic Outlook, April 2015.

Figure 3.3 plots the distribution of the *Consolidation* measure in my dataset. About 55% of the program/year cells include a condition on fiscal policy (Table 3.1). A positive number implies fiscal consolidation; a negative number fiscal expansion. The average for my dataset is +1.14% of GDP, implying an equivalent fiscal consolidation target (Table 3.1). However many programs actually require countries to undertake some degree of fiscal expansion with respect to the previous fiscal year. While this evidence might contradict popular understanding of IMF policies, there is evidence that spending might be increased during IMF programs, particularly social spending in developing countries (Nooruddin and Simmons [2006]; Independent Evaluation Office [2003]; Clements et.al [2013]). Also, governments might temporarily need to expand their fiscal capacity in times of crisis to prevent a collapse of economic activity (IEO [2003]).

My second measure of conditionality regards structural, or qualitative, conditions. This data also comes from the MONA database. I collect conditions in the realms of:

- (a) Public employment, wages, and civil sector reform;⁶⁴
- (b) Welfare system reform: social policy, pension, health, education, and social welfare reforms, including privatizations;⁶⁵
- (c) Labor market reforms, excluding reforms to the public sector, which are already included under category (a).⁶⁶

These are the types of conditions that are more likely to affect ordinary citizens' lives, and may face a stronger non-partisan opposition as a consequence.⁶⁷ According to the theoretical framework presented in Section 3.2, a relevant history of social unrest should weaken conditionality also in the qualitative policy areas defined by (a), (b) and (c). Labor market reforms, in particular, are negatively related to the relative strength of a country's labor unions (Caraway, Rickard and Anner [2012]).

My measure of structural conditionality is simple: I construct it as the yearly sum of the number conditions included in (a) (b) and (c) above:

$$SocialPolicyReform_{ct} = a_{ct} + b_{ct} + c_{ct} \tag{12}$$

64. MONA database, economic codes 7.4; 7.5; 7.6; 22

65. MONA database, economic codes 8; 8.1; 8.2;8.3;8.4;23;23.1;23.2

66. MONA database, economic codes 12 and 28.

67. Ponticelli and Voth [2011]; Tabellini and Passarelli[2013].

The average country undergoing an IMF program receives 1.5 conditions in *SocialPolicyReform* (Table 1); *Circa* 38% of the programs contain at least one condition in *SocialPolicyReform* (Figure 3.4) Measures of qualitative conditionality on international trade, or financial sector reforms are constructed in a similar way.⁶⁸ Conditions on international trade were included in 31% of the agreements, and conditions on financial sector reform were included in 22% of them (Figure 3.4).

Measuring structural conditionality is not a straightforward problem: the number of conditions counts, but scope and depth of the conditions also matter. A measure of conditionality based on the sum of the number of conditions is however a tradition in the literature on IMF conditionality: recent examples are Dreher and Jensen [2007], Woo [2013], or Gould [2003], who uses a three-year weighted average of the number of financial system reforms to evaluate the influence of private creditors. Stone [2008] uses a similar approach to evaluate the scope of IMF conditionality counting the number of policy domains covered by at least one condition.

Other authors have been trying to account for the relative stringency of conditionality. Caraway, Rickard and Anner [2012] assign different weights to a condition depending on whether this falls in the category of Prior Action, Structural Performance Criteria, or Structural Benchmark. This weighting procedure measures the relative importance of the condition from the point of view of the IMF (or how much the Fund cares about the condition being fulfilled), but it hardly provides a clear indication of the relative weight of the condition from a social or economic point of view. I therefore prefer a simple count variable, as this increases the transparency and interpretability of the results. However I also run some models using the weighting procedure proposed by Caraway, Rickard and Anner [2012] and show that my results are robust their measurement of structural conditionality.

3.3.3 Methodology and explanatory variables

To study the relationship between fiscal conditionality and social unrest, I rely on a panel data regression model of the type:

$$Y_{ct}=a+X_{ct-1}+AverageUnrest_{(t-1,t-2,t-3)}+\phi_c+\varepsilon_{ct} \quad (13)$$

68. As the sum of economic codes 1; 2; 26; 27 and 25; 25.1; 10; 10.1; 10.2; 10.3; 21.1; 21.2 respectively.

where Y_{ct} is the measure of conditionality (whether *Consolidation* or *SocialPolicyReform*) defined in Section 3.1. X_{ct-1} is a vector of country-year controls which include macroeconomic variables affecting fiscal consolidation, and several political and institutional factors which might affect fiscal conditionality, measured at $t-1$, in order to avoid a simultaneity bias. Country fixed effects (ϕ_c) aim at capturing time-invariant country-specific characteristics which might affect the degree of conditionality.

$AverageUnrest_{(t-1,t-2,t-3)}$ is the main effect of interest, which captures the short-term potential for social unrest in a given country. Following Ponticelli and Voth [2011], I construct this measure as a count variable of the episodes recorded in the Cross National Time Series Data Archive (CNTS).⁶⁹ The measure for social unrest ($AverageUnrest$) is computed as the three-year moving average of the number of episodes of social unrest occurring in country c prior to year t : this is, again, to avoid potential endogeneity arising from simultaneity and reverse causality. This is the dependent variable which I will use throughout the chapter to test the hypothesis that social unrest reduces the strength of conditionality in certain policy areas.⁷⁰

The CNTS dataset collects data on a broad range of social unrest episodes, taken from media reports of the incidents. These include general strikes, riots, anti government demonstrations, attempts to assassinate political figures, government crises, guerrilla warfare, revolutions and purges. However, $AverageUnrest$ only includes general strikes, riots and anti-government demonstrations. These variables potentially reflect different institutional structures, but all express forms of collective protest organized by the citizens/voters against the political establishment which might be triggered by undesired changes in economic policy. They also occur with relative frequency across the sample, while revolutions, for example, are extremely rare in this time period. Moreover, revolutions might easily be carried out by political forces other than citizens, such as the armed forces or opposition parties. As such, they don't necessarily reflect the citizens' capability and willingness to organize protests. A similar reasoning applies to guerrilla warfare. On the other hand, attempts to assassinate political figures may be the expressions of individual, rather than collective, grievance. Gov-

69. Banks [2013].

70. The average over three years is admittedly an arbitrary choice. This is however the same time span over which Gould [2003] builds a moving average of bank-friendly conditions. Three years is a short enough time span to reflect a country's recent history, and at the same time long enough to guarantee that a significant amount of country/year cells will have non-zero values.

ernment crises and purges, finally, reflect a completely different set of political and institutional dynamics, which run from a government to its citizens, rather than the other way around.

The objective nature of *AverageUnrest* forms an empirical basis for the common knowledge assumption at the basis of the theoretical model discussed in Section 3.2. Country officials might be willing to lie to the IMF about the potential for social unrest, but they cannot lie about the extent of what already happened and was recorded by the press. This variable is observed by the Fund in the same way as it is observed by the government.

As a methodological note, one common problem arises in the empirical literature on IMF lending: selection. The non-random selection of countries into IMF programs may affect the relationship between social unrest and austerity. In particular, a source of concern might be that governments self-select into IMF agreements in a way that is correlated with both austerity and social unrest. For example, a country could have entered two IMF agreements over a particular time span. Perhaps in the first occasion the government tried to fix its balance sheets in an attempt to avoid the IMF loan, while in the second occasion the government went directly to the negotiation stage with the Fund. It is likely that the pre-existing fiscal consolidation efforts reduced the need for fiscal adjustment in the first occasion versus the second (all else held equal). If the pre-program attempts to fiscal adjustment induced social unrest in the first episode, my model would capture a negative relationship between the planned fiscal conditionality and social unrest prior to the agreement. In order to exclude this possible source of bias, I run some robustness tests by including *actual* fiscal adjustment (in the year preceding the fiscal deficit condition) as an additional regressor. My results are robust to this specification.

Moreover, I estimate a selection model based on the Heckman procedure to explicitly control for selection. Implementing this procedure correctly proves challenging in this context, due to the unavailability of a credible instrument. Any variable related to the access to an IMF loan is very likely to be directly affecting the stringency of conditionality. This is because the political variables that have been used in the existing literature to this purpose (often in the context of program evaluation, such as the number of votes in line with the US in the UN general assembly) are not obviously exogenous.⁷¹ If a country can use its political influence to obtain more lending, it is also likely to use it to obtain a more

71. Barro, Lee [2005].

lenient deal, all else held equal. In fact, my results are indicative of an important influence of the Fund's major stakeholders in determining the stringency of conditionality, in line with the results of Dreher, Sturm and Vreeland [2009a,b].

In my selection equation, the exclusion restriction will be the presence of an IMF deal in the previous year. This variable has a strong predictive power on IMF program participation in any given year and therefore constitutes a relevant instrument.⁷² On the other hand, its validity as an instrument relies on the assumption that this variable affects the current degree of fiscal conditionality only to the extent to which prior IMF agreements (or their absence) induced the country to pursue prior fiscal consolidation. Therefore, conditional on pre-existing fiscal consolidation, IMF agreements taking place in the previous year should be unrelated to the degree of fiscal conditionality. My results hold to this specification, with similar magnitudes and significance as the baseline.⁷³

3.4 Results: fiscal conditionality and social unrest

This section discusses the results on quantitative fiscal targets and relative robustness tests. I also analyse how several political institutions might mediate between social unrest and fiscal conditionality.

3.4.1 Fiscal conditionality: macroeconomic determinants

The level of fiscal consolidations or expansion planned during an IMF program depends on a set of macroeconomic variables, such as the pre-existing level of fiscal imbalance of a country. Columns 1 and 2 of Table 3.2 displays the relationship between fiscal consolidation planned at the end of each program/year and a vector of macroeconomic and program-level controls, using a fixed-effects and a random-effects model, respectively.

In Column 1 the signs are generally as expected: the elasticity between the fiscal surplus prior to the agreement and *Consolidation* is negative and strongly significant, implying that an improvement in the fiscal balance worth one percentage point of GDP reduces fiscal consolidation targets by 0.11 percentage points of

72. Prior participation in IMF agreements has been used before as an instrument: see for example Rickard and Anner [2014].

73. Driscoll-Kraay [1998] standard errors are included whenever the dependent variable is continuous in order to take into account potential cross-sectional dependence. The error structure in this procedure is assumed to be correlated between the panels, autocorrelated up to two lags and robust to heteroskedasticity. I also run some robustness tests to prove that controlling parametrically for cross sectional dependence does not substantially change the results. Appendix A3.2 provides this set of estimations.

GDP. This means that one standard deviation increase in the budget balance reduces *Consolidation* by 16% of a standard deviation. The relationship between the trade balance and fiscal consolidation is also negative and significant, and has a similar magnitude: one standard deviation increase in the balance of payments reduces fiscal conditions by 21% of a standard deviation. Countries with a better balance of payments are required to carry out smaller budget cuts. One standard deviation increase in the debt-to-GDP ratio, on the other hand, leads to an increase in the fiscal consolidation target worth 52% of a standard deviation. The comparison of this magnitude with the one associated to budget balance (16%) suggests that fiscal consolidation targets depend primarily upon long-term fiscal sustainability concerns: a high debt/GDP ratio is more relevant in determining fiscal consolidation targets than a high deficit/GDP ratio.

The IMF's own growth forecast errors at the one-year horizon also display a positive relationship with fiscal consolidation targets (+0.14, significant at 5% level). This result can be interpreted in light of Blanchard and Leigh's [2013] findings. Under rational expectations, the correlation between growth forecast errors and planned fiscal adjustment should be zero. A positive coefficient is indicative of an underestimation from the IMF of the short-term multipliers arising from fiscal consolidation, which is consistent with Blanchard and Leigh's [2013] results.

Finally, the conditions can be targeted at the primary balance surplus, rather than the overall fiscal balance. Conditions on primary balance imply that the fiscal target should be calculated at the net of the payment of interests on outstanding public debt. As a major component of government expenditure is excluded, these conditions are nominally stricter: all else held equal, a condition on primary balance appears 6.8 percentage points of GDP higher than a condition on budget balance (significant at the 1% level).

Program-level variables do not seem to have much explanatory power on fiscal conditionality. The number of consecutive years a country has been under IMF arrangement has no significant effect on the dependent variable, nor does the overall IMF lending-to-GDP ratio. Also, the IMF lends to some low and middle-income member countries under concessional terms (interest rates on the loan are virtually zero and the repayment terms are longer). However, controlling for this factor does not seem to have a direct effect on fiscal consolidation targets. Column 2 displays the result of a random-effects model identical to the one in Column 1: the sign and magnitude of the effects are not dissimilar between the two models. However, the Hausman test rejects random-effects at the 1% confidence level,

suggesting that the fixed-effects model is preferable. All the estimations from this point onward will therefore include country fixed-effects.

3.4.2 Fiscal Conditionality and Social Unrest: Baseline

Table 3.3 introduces the social unrest measure, *AverageUnrest*: this is the moving average of the episodes of unrest (riots, anti-government demonstrations, and general strikes) which have occurred in the country during the three years prior to a program/year. Column 1 shows that *AverageUnrest* is negatively associated with *Consolidation*, and the magnitude of the effects is relevant: the elasticity is -0.62, implying that one standard deviation increase in this regressor (two episodes per year) reduces the dependent variable by 24% of a standard deviation (significant at the 1% level).

Columns 2-6 run some robustness checks on this baseline result. In particular, they control for possible sources of omitted variable bias. Previous literature on IMF conditionality has shown that several institutional and political factors, other than social unrest, influence the number and scope of policy prescriptions within IMF agreements. First of all, the number of protests might be related to the quality of democratic institutions: public protests might be less costly in democratic regimes than under autocratic governments. If transitioning to democracy simultaneously improves a country's budget balance, *AverageUnrest* would be measuring a spurious relationship between fiscal deficit targets and social unrest. In Column 2 I therefore include a control for the quality of democratic institutions, *Polity2*. This variable captures a regime's authoritarian spectrum on a 21-point scale ranging from -10 (hereditary monarchy) to +10 (consolidated democracy). While the level of democracy reduces the average fiscal consolidation target (elasticity of -0.18, significant at the 5% level), the coefficient associated with *AverageUnrest* is however largely unchanged, both in terms of magnitude and significance.

The degree of democracy prevalent in political institutions is not the only factor affecting fiscal conditionality and correlated with social unrest. The IMF might, for example, be reluctant to push costly reforms in election years, as citizens might punish the incumbent government in the polls. Starting negotiations with a new government once a program has already been agreed with the previous one is costly, and the Fund might try to avoid this scenario. At the same time, the years in the run-up to a general election might also be associated with more protests than usual. Column 3 tests whether timing of the elections matter for fiscal conditionality, taking into account the level of democracy. While the coef-

ficient on elections is negative (-0.25), it is not significant. Moreover, the effect of *AverageUnrest* is left largely unchanged (-0.61, significant at the 5% level).

Column 4 analyses another potential source of omitted variable bias: institutional, rather than extra-parliamentary, opposition. Within-parliament opposition might be related to extra-parliamentary opposition, as a reflection of the overall political fractionalization in society. However, the veto players internal to the parliament might be the true reason the IMF's deal is more lenient, as some reforms promoted by the IMF agreement might need to be voted on by the parliament. Therefore, when omitting a control for institutional fractionalization, the coefficient on *AverageUnrest* might be upwardly biased. However, Column 4 shows that veto players in the legislature do not seem to be driving the results. The number of veto players in a legislature is positively correlated with the stringency of fiscal conditionality (+0.79, significant at the 5% level). On the other hand, the effect of *AverageUnrest* is still negative and significant (-0.5, significant at the 5% level).

Another confounding factor might be the IMF's internal motives to be more or less indulgent towards a country. The literature on IMF programs and conditionality suggests that the Fund's most important contributors (the US and G5 nations) have not only influenced the IMF's decision to lend (Barro and Lee [2005]) but also the intensity and scope of conditionality attached to programs (Dreher Sturm and Vreeland [2009a,b]). The literature proxies for this geopolitical influence of the West on International Financial Institutions' operations using the number of votes a country casts in line with the US at the United Nations General Assembly (Barro and Lee [2005]), or with temporary membership of the UN Security Council (Dreher Sturm and Vreeland [2009a]). The external geopolitical pressure can reflect on the IMF's staff decision making process, to the point of making the Fund's win set larger when dealing with countries favoured by its powerful stakeholders (Dreher and Jensen [2007]).

To control for the IMF's win set, Column 5 includes both the number of votes in line with the US at the UN General Assembly as well as a dummy indicating whether a country is a temporary member of the UN Security Council in a given year. Consistently with Dreher Sturm and Vreeland [2009a,b], I find that political closeness to the US reduces the stringency of conditionality: one standard deviation increase in the votes in line with the US in any given year (0.12) reduces *Consolidation* by 0.84 percentage points of GDP, or 15% of its standard deviation. Temporary membership of the UN Security Council also reduces the need for

fiscal consolidation. The coefficient associated with *AverageUnrest* is however substantially unchanged from Column 1 and still strongly significant (1% level).⁷⁴

The models displayed in Table 3.3 include the full set of macroeconomic controls proposed in Table 3.2. This is in the attempt to control for the financial situation of a given country/year, which might drive fiscal consolidation efforts. However, the relationship between social unrest and fiscal conditionality might be driven by special economic circumstances which might not be fully captured by this set of macroeconomic controls. The need for fiscal consolidation might be more severe if a country borrows from the IMF to face a sovereign debt crisis, rather than a currency or banking crisis. At the same time, different types of financial crises could also carry different probabilities of social unrest, as they might have more or less direct effects on the citizens' welfare. In order to control for this potential source of omitted variable bias, I use the dataset compiled by Laeven and Valencia [2013], which measures the existence and typology of financial crises occurring in each country/year. Column 6 shows that while currency crises are associated with more lenient fiscal conditionality (as probably the IMF's intervention focuses on other policy areas in this case), the coefficient on *AverageUnrest* remains negative and strongly significant (-0.71, significant at the 1% level).

Overall, this section displays that the baseline results on the relationship between IMF fiscal conditionality and a recipient country's pre-existing history of social unrest are robust to a wide set of potential confounding factors.

3.4.3 Robustness to selection

Having excluded the most obvious candidates for a possible omitted variable bias in Table 3.3, Table 3.4 tests whether the relationship between social unrest and planned fiscal consolidation is robust to potential selection issues. This concern arises from the possibility of strategic behaviour on the side of the government in the run-up to an IMF program. If, for example, governments undertake fiscal adjustment when trying to stave off an IMF loan, they might generate social unrest in the process. If the attempt to avoid the IMF program fails, however, and the government subsequently requires the Fund to provide the country with lending assistance, my model would capture a negative relationship between planned fiscal consolidation efforts and social unrest prior to the agreement.

⁷⁴ Small differences in the coefficients are due to the loss of significant parts of the overall sample, when including multiple controls, which are not available for all program/years. Appendix A3.2 attempts to correct for this possible source of bias via interpolation of the additional control variables that are continuous in nature (*Polity2* and number of votes in line with the US at the UN General Assembly). The results are left substantially unchanged.

Table 3.4 therefore includes actual fiscal consolidation at year $t-1$ ($\Delta Deficit_{t-1} = Deficit_{t-1} - Deficit_{t-2}$) as an additional control. Column 1 shows that previous fiscal consolidation efforts have a strong negative effect on planned future fiscal consolidation: one percentage point improvement in the deficit-to-GDP ratio in previous years substantially reduces fiscal consolidation targets (coefficient 0.89, significant at the 1% level). However, the inclusion of this control does not considerably change the sign or magnitude of the coefficient associated with *AverageUnrest*: the elasticity is -0.52, significant at the 5% level. The magnitude of this effect is a reduction in fiscal targets worth 20% of a standard deviation for each standard deviation increase in *AverageUnrest* prior to the agreement: this is only a small decline from the magnitude of 24% estimated in Table 3.3.

Column 2 displays the effect of *AverageUnrest* only for the years subsequent to the first year of the program, to show whether the effect of social unrest differs at different stages of an IMF arrangement. The effect of social unrest seems to be irrelevant once the program has started: the IMF is not responsive to social unrest taking place once the government has agreed to a multi-year lending facility. However, Column 3 shows during the first year of a program, the effect is negative and strongly significant (1% level): one standard deviation increase in *AverageUnrest* prior to the beginning of a program reduces fiscal consolidation targets by 48% of its standard deviation in the first program/year. This, from a two-level game point of view, might suggest that the government reveals its options at the beginning of an IMF arrangement: the size of its win set is disclosed once and for all, and shifting domestic constraints might not be used against the Fund as a bargaining tool ex-post. The magnitude of this effect might however also suggest that a selection issue could be at play in the baseline results. In particular, countries more prone to social unrest might systematically be in need of less fiscal adjustment at the beginning of the agreement.

Column 4 is directly comparable to Column 1, as it includes all program years, not just the initial year. Here, however, I exclude Western European countries (Greece, Ireland, Portugal) from the estimation to control whether the effect is driven by this subset of IMF borrowers, who are characterised by higher levels of democratic accountability (as measured by the *Polity2* scale) with respect to the rest of the sample. Higher democratic accountability might imply more possibilities for the population to protest against the political authority. On the other hand, these countries might also be associated with a less severe need for fiscal adjustment (compared to other countries undergoing IMF arrangements). However, Column 4 shows that this is not the case: the coefficient on *AverageUn-*

rest is substantially unchanged from Column 1: Western European democracies are not driving the results.

However, these results do not rule out the possibility of a selection bias. Column 5 therefore presents the estimation of the Heckman selection model, in further attempt to correct for selection into IMF programs. The selection equation (Column 6) includes the same vector of controls as in the outcome equation, and an exclusion restriction: this is a dummy indicating whether the country was under an IMF program at year $t-1$. While this variable has strong predictive capacity for future inclusion into an IMF lending program (+1.13, significant at the 1% level), it should have no direct or indirect effect on *Consolidation*, once controlling for fiscal reform that had already taken place in the same country prior to the current program/year. Column 5 shows that this further control for selection leaves the magnitude and significance of the coefficient associated with *AverageUnrest* substantially unchanged from Column 1. Practically, this implies that two additional episodes of civil unrest in each of the three years prior to any program/year reduce fiscal consolidation targets by 1.2 percentage points of the deficit-to-GDP ratio (21% of a standard deviation).

3.4.4 Mediating factors: social unrest and domestic politics

The previous sections have provided evidence of the negative relationship between fiscal consolidation under IMF arrangements and the level of extra-parliamentary constraints faced by a government. In particular, if a population is generally more prone to protests in the run-up to an IMF program review, its governments receives less stringent fiscal consolidation conditions in the following program review.

However, the strength of this constraint is unlikely to be the same for all countries, because different governments face different incentives to respond to their citizens' needs and demands. Leaders in fully democratic countries, for example, might be more responsive to the threat of social unrest than more autocratic leaders. Also, the electoral cycle might matter: being close to the elections might make politicians reluctant to entertain the idea of costly reforms that could trigger protests. Partisanship might as well play an important role. A party that witnesses its electoral base being largely represented among the protesters is probably more likely to be responsive to their needs. Finally, extra-parliamentary opposition might interact with intra-legislature fragmentation: if governments are constrained by the protests and also face a large number of veto players in the parliament, the effect of social unrest on fiscal targets might be enhanced.

Table 3.5 addresses how these political variables mediate between social unrest and fiscal conditionality by exploiting a set of interaction terms with political variables taken from the World Bank Database of Political Institutions. Column 1 shows that social unrest has a stronger effect on *Consolidation* the higher the degree of democratic accountability of a country's institutions. The coefficient on the interaction term between *AverageUnrest* and the polity score is negative (-0.25, significant at the 5% level), implying that the higher the level of democratic accountability, the larger the effect of *AverageUnrest* on fiscal conditionality. This can be read as evidence that autocratic governments might therefore be unwilling or unable to use this particular type of domestic constraint as a bargaining tool when negotiating with the Fund. On the other hand, the extent of intra-parliamentary opposition seems to have no mediating effect between social unrest and fiscal consolidation: in Column 2 the interaction term between the number of veto players in a legislature and social unrest prior to the agreement is not significant.

Column 3 shows that election timing matters: social unrest has a stronger negative effect on fiscal consolidation targets during electoral years (-0.38, significant at the 10% level). This result, together with the evidence about the differential role of social unrest along the *Polity2* scale, suggests that the electoral cycle might be an important key to an understanding of the system of incentives faced by a government when bargaining over conditionality with the IMF. The higher the level of democratic accountability, and the closer in time the elections are, the more the need of a government to be responsive to its electoral base. Therefore, the larger the potential for extra-parliamentary opposition to provide leverage in the negotiation process, as the threat of involuntary defection becomes more credible.

Finally, the results in Columns 4 and 5 show how political leanings of the government interact with the measure of social unrest. Column 4 shows that those government classified as right-wing by the World Bank Database of Political Institutions do not use social unrest as an instrument to obtain fewer conditions. In fact, they might be using it to obtain *stricter* fiscal conditionality, as the interaction term between social unrest and fiscal consolidation is actually positive (+0.42, significant at the 10% level). The opposite is true for left-wing governments: when these are faced with higher *AverageUnrest*, they negotiate less stringent fiscal consolidation targets (-1.48 percentage points, significant at the 10% level).

These two results can be interpreted as further evidence that the citizen/voter is pivotal in understanding how social unrest can affect fiscal conditions. Left-wing governments are more likely to find institutional and political support among members of the trade unions and social movements. Therefore, since their electoral base is likely to be disproportionately represented in strikes and anti-government demonstrations, the threat of involuntary defection can become severe for these governments, unless they accommodate some of their voters' demands. This view is in line with Rickard and Anner's [2014] results. On the other hand, conservative governments do not typically have their electoral base in the trade unions or social movements. As their electoral support finds its root in other strata of society, the threat of involuntary defection due to social unrest becomes less credible for these governments. Moreover, conservative politicians might even be willing to use the occasion of an IMF agreement to push reforms that might have been previously hardly attainable due to domestic opposition, as the positive sign on the interaction term in Column 4 seems to suggest.

3.5 Social Unrest and structural conditions

Citizens may care about the size of a budget cut, but they are also likely to care about its distribution. If a government is constrained when implementing fiscal adjustment because of the threat of social unrest, it should also have difficulties implementing structural changes in labour markets, welfare state and public employment, as these policy realms are among those that will impact citizens more directly. Therefore, according to the two-level game theoretical framework, the evidence of strong social movements should also reduce planned structural conditionality in the economic policy domains.

3.5.1 Results

The dependent variable in Table 3.6 is the number of conditions in *SocialPolicyReform* in each program/year. Since the dependent variable is discrete and presents overdispersion I use a negative binomial regression framework.⁷⁵

Column 1 shows that the elasticity between *AverageUnrest* and the number of conditions in *SocialPolicyReform* is negative (-0.12) and statistically significant at the 5% level: one standard deviation increase in the measure of social unrest reduces the number of these conditions by 0.25 (a quarter of a condition), or 11% of a standard deviation.

⁷⁵. The standard deviation of *SocialPolicyReform* is about 1.5 times larger than its mean (Table 3.1).

Column 2 tests whether the result is robust to a classification of structural conditions which takes into account their relevance, using the weighting system proposed by Caraway Rickard and Anner [2012]. In this categorization, Prior Actions receive a weight of 4, Structural and Quantitative Performance Criteria a weight of 3, and Structural Benchmarks of Indicative Targets a weight of 2. Column 2 shows that introducing this additional classification of the type of conditions does not change the direction or magnitude of the baseline result. However, while the coefficient is similar to Column 1, the higher standard deviation associated with this weighted dependent variable reduces the magnitude of the results substantially: one standard deviation increase in *AverageUnrest* is now associated with a decline in the weighted number of reforms worth 4% of its standard deviation, about a fourth of the magnitude found using the unweighted structural conditionality measure.

Columns 3 and 4 run a placebo test: the dependent variables are the number of conditions associated with trade reform and financial sector reform in any given program/year. These are also conditions often included in IMF programs and they might face some form of domestic opposition. However, the type of constraints associated with reforms of these policy areas are less likely to depend upon protests and more likely to hinge on the influence of business associations, for example, or other institutional constraints (Woo [2013]). Gould [2003, 2006] shows that financial sector conditions, in particular, are shaped by the presence and influence of commercial banks among the group of lenders working in partnership with the IMF. Column 3 shows that there is no significant relationship between the number of trade conditions and social unrest prior to a program/year. Column 4 shows that the number of financial sector reforms is actually moderately increasing in the number of episodes of unrest prior to a program/year. This is evidence that the decline in the number of conditions following episodes of social unrest is not simply reflecting a generalized reduction in conditionality. Rather, it specifically reflects a decline in the number of conditions related to public expenditure.

Finally, Column 5 adds controls for whether the country is undergoing a banking, currency, or sovereign debt crisis, as classified by Laeven and Valencia [2013]. Banking crises are associated with fewer social policy conditions (-1.5, significant at the 1% level), while the other two types of crises do not have a significant impact on the number of conditions. However, adding these controls does not change the sign or magnitude of *AverageUnrest* substantially: the estimated elasticity is -0.1, significant at the 10% level.

Table 3.7 displays the same results as in Table 3.6 using a zero inflated negative binomial model, as a robustness test. This model should be able to take into account excess zeros obtained by the data generating process.⁷⁶ Column 1 shows that the zero-inflated model has to be preferred to the standard negative binomial MLE: the Z statistic of a significant difference between the two models is positive (+1.76) and statistically significant at the 5% level. Column 1 also largely confirms that one standard deviation increase in the episodes of social unrest reduces the number of conditions on social policy reform by 12% of a standard deviation (significant at the 5% level). Columns 2-5 test the robustness of this result to the inclusion of additional controls.

Column 2 adds controls for the macroeconomic determinants of fiscal consolidation, proposed in Table 3.2, plus a dummy variable indicating whether the program also features a condition on fiscal deficit in the same year. The presence of a condition on fiscal deficit actually reduces the number of structural conditions, suggesting that these two types of conditions are not complementary, but rather alternative. Nevertheless, the coefficient associated to *AverageUnrest* remains negative (-0.12) and statistically significant at the 10% level.

Column 3 adds controls for other domestic political constraints, explored in Table 3.2.⁷⁷ The level of democratic accountability has no significant effect on the number of *SocialPolicyReform* conditions. The presence of elections in the same program/year does not affect the dependent variable in a significant fashion either. However, more veto players in the legislature are associated with more stringent structural conditionality, as it is the case with fiscal conditionality (+0.28, significant at the 5% level).

Structural conditionality may also be driven by the type of economic crisis the country is facing. A sovereign debt crisis, for example, might suggest more conditions targeted at reducing the public expenditure bill. Column 4 therefore adds controls for the type of crisis, as coded in Laven and Valencia's [2013] dataset. While the coefficient on banking crises is negative and strongly significant, the other two types of crises do not have any statistically significant impact on the dependent variable. Moreover, the coefficient on *AverageUnrest* is substantially unchanged (-0.11) and remains significant at the 5% level.

⁷⁶. Many programs do not include any conditions on *SocialPolicyReform* at all, possibly because the type of crises these countries are facing do not reflect a fiscal imbalance.

⁷⁷. I run separate models in this case because the ZINB procedure fails to converge when introducing too many regressors.

Finally, Column 5 tests for the influence of geopolitical influences which might affect the IMF's win-set. I include controls for the number of votes in line with the US at the UN General Assembly and a dummy for whether the country is a temporary member of the UN Security Council. The effect of voting at the UN, while negative, is not statistically significant. On the other hand membership of the Security Council increases the number of structural conditions. The coefficient associated with *AverageUnrest* remains negative (-0.11) and significant at the 5% level.

The findings in Tables 3.6 and 3.7 suggest a negative relationship between social policy conditions and *AverageUnrest*. This provides further evidence in support of the domestic constraints hypothesis: if governments are more constrained by strong social movements, they are asked to implement fewer structural conditions in the realms that relate most closely to public sector reform.

3.6 Conclusions

This chapter analyses how a country's propensity for social unrest influences fiscal policy choices in the context of IMF arrangements. To study the relationship between extra-parliamentary opposition and fiscal conditionality, I construct an original dataset on fiscal policy and government budget reforms included in all IMF agreements between 2002 and 2012.

If citizens of a country are more prone to social unrest, the government might postpone fiscal stabilisation and other social policy-related reforms in order to avoid the welfare loss resulting from rioting (Ponticelli, Voth [2011]; Tabellini Passarelli [2013]). If this is a credible threat, however, it might also influence the negotiations between a country and the IMF, by providing the government with a bargaining advantage towards the Fund. The theoretical basis is Putnam's [1988] seminal contribution on two-level games. This theory predicts that if a party is constrained by domestic constituencies, and this constraint is common knowledge between the negotiating parties, then the most constrained negotiator obtains a better deal, because of the threat of involuntary defection. As a result, conditionality in the realms of fiscal and social policy should be fewer, and less stringent.

My findings strongly support this hypothesis: an increase in social unrest prior to a program/year significantly reduces fiscal consolidation targets. Moreover, governments also obtain fewer structural conditions related to welfare provision and public employment. These results are robust to controlling for a vast array

of variables deemed relevant by previous literature on IMF conditionality, and to potential selection issues.

The analysis presented here might benefit from a measurement of the scope and intrusiveness of structural conditions. Further work is also required in order to understand the degree of compliance with fiscal conditions. The simple model presented here predicts that compliance with fiscal conditionality should be *quasi-perfect*, given that domestic constraints are taken into account as a result of first-stage negotiations. As this is not likely to be the case, other variables may intervene in making defection (whether voluntary or not) arise in equilibrium. Finally, it would be interesting to study whether social unrest influences governments' decisions to undertake fiscal consolidation outside the context of IMF programs.

This chapter shows that social movements are able to influence economic policy-making, even in the context of IMF lending. This is an interesting result, since it suggests that extra-parliamentary opposition might be an additional political force shaping fiscal policy choices, alongside the formal institutions already considered by the political economics literature. A study of the domestic political constraints to the implementation of economic policy, moreover, might help explain why a similar set of policies might produce different effects in different contexts. Fiscal reforms typically have many social and distributional consequences, and optimal fiscal policy design may substantially differ from the textbook approach, once these social constraints are taken into account.

IMF programs might be considered a peculiar case to study the effects of political constituencies on fiscal policy, but it is precisely in virtue of this peculiarity that they constitute an informative example. If the threat of civil unrest is capable of restricting options even under IMF arrangements, then it may also explain fiscal policy choices in more conventional times.

Figures

Figure 3.1: An increase in the domestic constraints shifts the set of possible agreements from $[Sm1, G^*]$ to $[Sm2, G^*]$.




Figure 3.2. Countries undergoing IMF lending programs, 2002-2012. Source: IMF MONA Database




Figure 3.3: Distribution of fiscal consolidation measure as defined in Equation n.11. A positive number imply fiscal consolidation, negative ones fiscal expansions. Years 2002-2012 Source: IMF MONA data-base.





Figure 3.4: Distribution of structural conditionality measures. Number of conditions per program/year, as defined by Equation n12. Years 2002-2012. Source: IMF MONA Database.



Tables

Table 3.1: Descriptive statistics. Sources: IMF Monitoring of Fund Arrangements Database, Cross National Time Series Data Archive, World Economic Outlook (2015).

Variable	Obs	Mean	Std. Dev.	Min	Max
Program Characteristics					
#Consecutive Years In Program	486	3.33	2.61	0.00	13.00
Concessional Lending	483	0.56	0.50	0.00	1.00
Deficit Condition	486	0.56	0.50	0.00	1.00
Access/GDP	470	5.59	18.18	0.00	138.74
Conditionality					
Consolidation	231	1.11	5.61	-17.52	19.57
Social Policy Reform, Yearly #	472	1.27	2.32	0.00	16.00
Trade Reform, Yearly #	472	1.10	2.32	0.00	23.00
Financial Sector Reform, Yearly #	472	0.59	1.52	0.00	10.00
Controls: Social Unrest					
Average Unrest (three year moving average)	480	0.69	2.15	0.00	28.33
General Strikes, Yearly #	485	0.09	0.58	0.00	4.00
Riots Yearly #	485	0.25	0.80	0.00	5.00
Anti-Government Demonstrations, Yearly #	485	0.48	1.51	0.00	8.00
Macroeconomic Variables					
Average Forecast Error GDP Growth (t+1)	481	0.43	4.48	-20.43	40.25
Fiscal Deficit/GDP	459	-2.50	8.13	-23.48	125.45
Balance of Payment/GDP	472	-7.42	8.79	-70.98	19.75
Public Debt/GDP	429	71.92	80.64	9.55	931.64

Table 3.2: Fiscal conditionality: macroeconomic and program-related controls

VARIABLES	(1) Consolidation FE	(2) Consolidation RE
Forecast Error Growth (One year)	0.140** (0.057)	0.108* (0.057)
Deficit t-2	-0.112*** (0.026)	-0.102*** (0.025)
BOP t-1	-0.138*** (0.045)	-0.138*** (0.041)
Debt t-1	0.037*** (0.008)	0.037*** (0.007)
Primary Balance	6.842*** (1.741)	2.336** (1.021)
Years in Program	-0.175 (0.185)	-0.124 (0.167)
Concessional	2.179 (1.533)	0.008 (1.003)
Access/GDP	-0.120 (0.073)	-0.028 (0.021)
Constant	-4.656*** (1.444)	-3.336*** (1.128)
Observations	202	202
Within R-squared	0.325	0.277
Number of Countries	61	61
Hausman Test χ^2	21.43	
P Value	0.006	

OLS regression. Sources: MONA database and World Economic Outlook (April 2015). The dependent variable is fiscal consolidation at time t (only country/year cells that received a condition on fiscal balance are included). Controls include: the forecast error of GDP growth (one year horizon); Government deficit/GDP at $t-2$; Balance of Payments/GDP and Central Government Debt/GDP at $t-1$; a dummy for whether the condition refers to government primary balance; number of consecutive years under an IMF program; a dummy for whether the IMF program is classified as concessional lending or not; Access (in SDR) as a proportion of current national GDP. Robust standard errors in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 3.3. Fiscal Consolidation and Social Unrest: Baseline

VARIABLES	(1) Consolidation	(2) Consolidation	(3) Consolidation	(4) Consolidation	(5) Consolidation	(6) Consolidation
AverageUnrest	-0.652*** (0.189)	-0.630*** (0.176)	-0.615** (0.200)	-0.506** (0.168)	-0.636** (0.225)	-0.716*** (0.210)
Polity2		-0.182** (0.077)	-0.184 (0.102)			
Election Years			-0.254 (0.139)			
Veto Players				0.797** (0.248)		
Member UNSC					-4.005** (1.240)	
Votes with US					-7.080*** (1.975)	
Banking Crisis						-1.549 (1.139)
Currency Crisis						-2.469* (1.109)
Sov. Debt Crisis						-0.885 (1.736)
Controls Table 2	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Constant	-3.737** (1.296)	1.950 (1.457)	1.355 (2.020)	-2.890 (1.899)	0.275 (0.869)	-3.209** (1.233)
Observations	188	164	160	168	151	183
Number of groups	59	52	51	54	56	57

OLS regression. The dependent variable is fiscal conditionality at time t . Column (1) includes all program year/cells which have a condition on fiscal deficit balance. Column (3) adds controls for the level of democratic accountability in a given country/year, measured by the polity2 scale; Column (4) includes also controls for whether the country is in an election year (WB Database of Political Institutions). Column (5) controls for the number of veto players in a legislature. Column (6) adds controls for geopolitical influences, namely: votes in line with the US at the UN General Assembly and whether the country is a temporary member of the UN Security Council. Column (6) controls for the type of economic crisis. All Columns include the full set of controls included in Table 2, as well as country fixed effects. Driscoll-Kraay standard errors in parentheses.*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.4. Fiscal Consolidation and Social Unrest: Robustness to Selection

VARIABLES	(1)		(2)		(3)		(4)		(5)		(6)	
	Consolidation		Consolidation		Consolidation		Consolidation		Consolidation		Program Heckman Selection Probit	
	OLS		OLS		OLS		OLS		OLS		OLS	
AverageUnrest	-0.522** (0.227)		0.113 (0.121)		-1.278*** (0.316)		-0.528** (0.233)		-0.561** (0.191)		-0.052 (0.042)	
Consolid. t-1	-0.887*** (0.041)		-0.864*** (0.072)		-0.774*** (0.106)		-0.887*** (0.041)		-0.888*** (0.074)		-0.035** (0.017)	
Program t-1											1.353*** (0.120)	
Lambda											-1.161*** (0.268)	
Controls Table 2												
Country FE	Yes		Yes		Yes		Yes		Yes		Yes	Yes
	Yes		Yes		Yes		Yes		Yes		Yes	Yes
Constant	-2.985*** (0.509)		-2.643* (1.326)		-3.959** (1.260)		-3.179*** (0.499)		-1.277** (0.518)		1.557 (1.113)	
Observations	188		107		81		180		165		670	
Number of Country	59		53		55		56		55			

In Cols. (1)-(5) the dependent variable is planned fiscal consolidation at the end of year t , as recorded by the MONA database. In Column (6) is a dummy indicating whether the country is under an IMF program in year t . All columns controls for past fiscal consolidation (measured as the change in the actual fiscal balance between $t-2$ and $t-1$). Column (1) includes all country/years; Column (2) only the program years after the first; Column (3) the first year of every program. Column (4) excludes Ireland, Greece and Portugal. n. Column (5) is the outcome equation of the Heckman model, which includes the inverse Mills ratio resulting from Column (6). Column (6) is a probit model for selection into an IMF program, which includes past participation in an IMF program as the exclusion restriction. All Columns include the macroeconomic controls included in Table (1). Driscoll Kraay standard errors in parentheses in Cols (1)-(5) and robust standard errors in Column (6). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 3.5: Fiscal consolidation and social unrest: interaction with institutional factors

VARIABLES	(1) Consolidation	(2) Consolidation	(3) Consolidation	(4) Consolidation	(5) Consolidation
Unrest*Polity	-0.258** (0.107)				
Polity	-0.147 (0.088)				
Unrest*VetoPl.		0.015 (0.184)			
Veto Players		0.771** (0.260)			
Unrest*Elections			-0.383* (0.180)		
Elections			0.145 (0.178)		
Unrest*Conservative				0.425* (0.196)	
Conservative Govt.				-2.356** (0.918)	
Unrest*Left					-1.485* (0.729)
Left					1.555 (0.885)
Average Unrest	1.356 (0.782)	-0.560 (0.638)	-0.330 (0.211)	-0.797*** (0.205)	-0.329* (0.172)
Controls Table 1	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes
Constant	1.428 (1.505)	-2.761 (2.198)	0.047 (1.175)	-3.456** (1.455)	-4.270*** (1.228)
Observations	164	168	168	188	188
Number of groups	52	54	54	59	59

OLS regression. The dependent variable is planned fiscal consolidation in year t . This table runs a series of models where average unrest is interacted with a set of controls for domestic politics from the World Bank database of political institutions. Column (1) controls for the interaction between average unrest and polity score. Column (2) for veto players in a legislature, measured as the number of checks and balances. Col.(3) for the interaction between unrest and election years. Cols. (4) and (5) for the interaction between unrest and right and left wing governments, respectively. All Columns include the full set of controls included in Table 1 as well as country fixed effects. Driscoll-Kraay standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.6: Structural conditionality: negative binomial regression

VARIABLES	(1) Social Policy Reform	(2) Social Policy Reform (CRA weights)	(3) Trade Policy Reform	(4) Financial Sector Reform	(5) Social Policy Reform
AverageUnrest	-0.120** (0.051)	-0.128* (0.070)	-0.074 (0.122)	0.117* (0.061)	-0.103* (0.057)
AccessGDP	0.016 (0.025)	0.028 (0.047)	-0.015 (0.035)	-0.038 (0.178)	0.017 (0.025)
Banking Crisis					-1.520*** (0.339)
Currency Crisis					0.307 (0.609)
Sov. Debt Crisis					-1.076 (1.223)
Country FE	Yes	Yes	Yes	Yes	Yes
Constant	0.632*** (0.164)	1.770*** (0.240)	0.894** (0.450)	-1.611*** (0.188)	0.379*** (0.134)
Observations	431	431	431	431	412

Negative binomial regression. The dependent variable in Column (1) is the number of conditions in social policy reform recorded at year t . In Column (2) is the number of conditions weighted according to their nature, following the weighting procedure proposed by Caraway, Rickard and Anner (2012). In Cols. (3)-(4) is the number of conditions on international trade reform and financial sector reform, respectively. Column (5) adds controls for the type of crisis, according to Laeven and Valencia's (2012) classification. Country fixed effects (unconditional). Robust standard errors in parentheses.*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.7: Structural conditionality: zero-inflated negative binomial

VARIABLES	(1) Social Policy Reform	(2) Social Policy Reform	(3) Social Policy Reform	(4) Social Policy Reform	(5) Social Policy Reform
Average Unrest	-0.131** (0.052)	-0.125* (0.065)	-0.113*** (0.043)	-0.118** (0.057)	-0.116** (0.050)
Access/GDP	-0.004 (0.034)	-0.027* (0.016)	-0.044*** (0.011)	-0.004 (0.031)	0.001 (0.027)
Debt t-1		-0.005 (0.003)			
BOP t-1		-0.009 (0.011)			
Deficit t-2		0.030 (0.024)			
Fiscal Consolidation		-0.664*** (0.243)			
Polity			0.015 (0.155)		
Veto Players			0.283** (0.121)		
Election			-0.117 (0.101)		
Banking Crisis				-1.441*** (0.261)	
Currency Crisis				0.057 (0.536)	
Sov. Debt Crisis				-0.019 (0.384)	
Membership UNSC					0.904*** (0.284)
Votes with US (UNGA)					-0.478 (1.780)
Country FE	Yes	Yes	Yes	Yes	Yes
Constant	0.867*** (0.160)	1.582*** (0.261)	0.674 (1.404)	1.292*** (0.179)	0.907** (0.407)
Vuong Test	1.76**				
Observations	431	375	373	412	364

Zero-inflated negative binomial regression. The dependent variable in cols. (1)-(5) is the count of conditions in social policy reform in a given country/year. Column (2) adds macroeconomic controls. Column (3) controls for the level of democratic accountability (polity scale), the number of veto players in a legislature and whether the program/year is also an election year. Column (4) controls for the type of crisis, as in Laeven and Valencia (2012). Column (5) controls for geopolitical influences (votes in line with the US at the UN General Assembly and whether the country is a temporary member of the UN Security Council). Unconditional country fixed-effects are included in all models. Robust standard errors in parentheses.*** p<0.01, ** p<0.05, * p<0.1.

Appendix A3.1: Cross sectional dependence

Cross sectional dependence is a common problem arising in panel datasets where units are not randomly sampled. Countries, for example, are subject to regional shocks and a macroeconomic crisis can expand rapidly across borders, generating regional (or even global) common disturbances. Assuming the residuals are correlated within groups, but uncorrelated across groups, might lead to underestimate the standard errors (Driscoll and Kraay [1998]). Corrections for spatial dependence in macro datasets are possible, but they generally rely on the assumption of a large time dimension ($T > N$), which is unrealistic in many applications (Hoeckle [2007]). Controlling parametrically for cross-sectional dependence, on the other hand, imposes a precise structure on the data and could be prone to misspecification. Driscoll and Kraay's [1998] procedure allows one to non-parametrically correct for cross-sectional dependence, autocorrelation, and heteroskedasticity, while at the same time not requiring a large T/N ratio. Hoeckle [2007] performs the Monte Carlo simulations on a sample where $N=1000$ and the maximum $T=40$ ($T/N=0.04$). The sample I use in this chapter includes 67 countries undergoing fiscal conditionality, with a maximum number of years in which fiscal conditions are recorded in any given country being eight ($T/N=0.12$). This technique is overall more flexible than controlling parametrically for cross sectional dependence. Nevertheless, Table A.3.2 shows that including year dummies with robust standard errors does not substantially affect the results, compared to Table 3.3.

Table A3.2: Controlling parametrically for cross-sectional dependence

VARIABLES	(1) Consolidation	(2) Consolidation	(3) Consolidation	(4) Consolidation	(5) Consolidation
AverageUnrest	-1.068** (0.468)	-0.977** (0.434)	-1.022** (0.458)	-0.675** (0.328)	-1.136** (0.493)
Forec.error (1Y)	0.170*** (0.060)	0.075 (0.048)	0.069 (0.050)	0.175*** (0.049)	0.145*** (0.049)
Deficit t-2	-0.099** (0.048)	-0.142*** (0.035)	-0.147*** (0.032)	0.153*** (0.047)	-0.106** (0.050)
BOP t-1	-0.133* (0.073)	-0.115 (0.073)	-0.135* (0.076)	-0.082 (0.058)	-0.130* (0.072)
Debt t-1	0.037*** (0.010)	0.001 (0.021)	0.005 (0.023)	0.019* (0.011)	0.037*** (0.011)
Primary Balance	7.350*** (1.782)	7.464*** (1.957)	7.418*** (1.946)	6.788*** (1.721)	7.180*** (1.729)
Years in program	-0.216 (0.260)	-0.399 (0.302)	-0.426 (0.302)	-0.420 (0.252)	-0.246 (0.276)
Concessional	1.405 (2.131)	1.965 (2.757)	2.113 (2.745)	2.062 (1.425)	1.472 (2.109)
Access/GDP	-0.198** (0.076)	-0.223** (0.088)	-0.225** (0.089)	-0.285*** (0.055)	-0.182** (0.083)
Election			0.020 (0.363)		
Polity2		-0.115 (0.186)	-0.146 (0.228)		
Membership UNSC				-6.228*** (1.791)	
Votes with US (UNGA)				-5.299 (5.877)	
Banking Crisis					-1.269 (1.926)
Currency Crisis					-3.612 (2.163)
Sov.Debt Crisis					-0.182 (1.673)
Year FE	Yes	Yes	Yes	Yes	Yes
Constant	-4.533** (2.121)	-0.156 (2.900)	-0.596 (3.207)	0.082 (1.698)	-3.773* (2.221)
Observations	188	164	160	151	183
R-squared	0.378	0.430	0.462	0.417	0.413
Number of Countries	59	52	51	56	57

OLS regression. The dependent variable is planned fiscal consolidation at the end of year t. This Table reposes the model in Table 2, but using a different simple robust standard errors, instead of Driscoll-Kraay standard errors. I control for cross-sectional dependence by adding year fixed effects. Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Appendix A3.2: Interpolation of dependent variables

Table A3.1. Interpolation of political variables

VARIABLES	(1) Consolidation	(2) Consolidation
AverageUnrest	-0.585*** (0.148)	-0.617*** (0.184)
Forecast Err.Growth	0.112** (0.035)	0.148*** (0.033)
Deficit t-2	-0.120** (0.038)	-0.097* (0.048)
BOP t-1	-0.131*** (0.024)	-0.145*** (0.025)
Debt t-1	0.005 (0.015)	0.037*** (0.004)
Primary Balance	5.047*** (0.620)	5.247*** (0.377)
Polity2 (Interpolation)	-0.105 (0.089)	
Votes with US (Interp)		1.354 (2.448)
Country FE	Yes	Yes
Constant	-1.541* (0.775)	-4.860*** (0.890)
Observations	166	188
Number of groups	53	59

OLS regression. The dependent variable is planned fiscal consolidation at the end of year t . Column (1) controls for the effect of the Polity scale once the missing cells for this variables have been replace by their country-level average throughout the years. Column (2) does the same with the number of votes in line with the US at the UN General Assembly. Driscoll-Kraay standard errors in parentheses.*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

References

- Adelino, M., Schoar, A., & Severino, F. [2016]. Loan originations and defaults in the mortgage crisis: The role of the middle class. *Review of Financial Studies*, hhw018.
- Alesina, A., & Ardagna, S. [2010]. Large changes in fiscal policy: taxes versus spending. In *Tax Policy and the Economy*, Volume 24 [pp. 35-68]. The University of Chicago Press.
- Alesina, A., & Drazen, A. [1991]. Why Are Stabilizations Delayed?. *The American Economic Review*, 81[5], 1170-1188.
- Alesina, A., & Passalacqua, A. [2015]. The Political Economy of Government Debt . *National Bureau of Economic Research*. [No. w21821]
- Alesina, A., & Perotti, R. [1995]. The political economy of budget deficits. *Staff Papers*, 42(1), 1-31.
- Alesina, A., Carloni, D., & Lecce, G. [2012]. The electoral consequences of large fiscal adjustments. In *Fiscal Policy after the Financial Crisis* [pp. 531-570]. University of Chicago Press.
- Angrist, J. D., & Pischke, J. S. [2008]. Mostly harmless econometrics: An empiricist's companion. *Princeton University Press*.
- Aron, J., Duca, J. V., Muellbauer, J., Murata, K., & Murphy, A. [2012]. Credit, housing collateral, and consumption: evidence from Japan, the UK, and the US. *Review of Income and Wealth*, 58(3), 397-423.
- Bailey, M., Cao, R., Kuchler, T., & Stroebel, J. [2016]. *Social Networks and Housing Markets National Bureau of Economic Research* (No. w22258).
- Banks A. Wilson K.A., [2013]. Cross National Time Series Data Archive. *Databanks Int. Jerusalem, Israel*.
- Barberis N., Shleifer A., Vishny R. [1998]. A model of investor sentiment. *Journal of Financial Economics* 49.3 (1998): 307-343.
- Barro, R. J., & Lee, J. W. [2005]. IMF programs: Who is chosen and what are the effects?. *Journal of Monetary Economics*, 52[7], 1245-1269.
- Bayer, P., Mangum, K., & Roberts, J. W. [2016]. Speculative Fever: Investor Contagion in the Housing Bubble. *National Bureau of Economic Research* (No. w22065).

- Bertrand, M., Duflo E. & Mullainathan S. [2004]. How Much Should We Trust Differences-In-Differences Estimates?. *The Quarterly Journal of Economics* 119, 249– 275.
- Bertrand, M., & Morse, A. [2016]. Trickle-Down Consumption. *Review of Economics and Statistics* 2016 98:5, 863-879
- Blanchard, O. J., & Leigh, D. [2013]. Growth Forecast Errors and Fiscal Multipliers. *The American Economic Review*, 103[3], 117-120.
- Blundell, R., Pistaferri, L., & Preston, I. [2008]. Consumption inequality and partial insurance. *The American Economic Review*, 1887-1921.
- Bordalo, P., Gennaioli, N., & Shleifer, A. [2016]. Diagnostic Expectations and Credit Cycles *National Bureau of Economic Research*. (No. w22266)
- Bover, O. (2015). Measuring expectations from household surveys: new results on subjective probabilities of future house prices. *SERIEs*, 6(4), 361-405.
- Bowles, S., & Park, Y. [2005]. Emulation, Inequality, and Work Hours: Was Thorsten Veblen Right?. *The Economic Journal*, 115(507).
- Branch, W. A. (2004). The theory of rationally heterogeneous expectations: evidence from survey data on inflation expectations. *The Economic Journal*, 114(497), 592-621.
- Brunnermeier, Markus K, and Martin Oehmke [2013]. Bubbles, Financial Crises, and Systemic Risk. *Handbook of the Economics of Finance*. Amsterdam: Elsevier
- Burnside C., Eichenbaum M., Rebelo S. [2014]. Understanding booms and busts in housing markets. Working Paper.
- Burke, M. A., & Manz, M. (2014). Economic literacy and inflation expectations: evidence from a laboratory experiment. *Journal of Money, Credit and Banking*, 46(7), 1421-1456.
- Caballero R., and Simsek A. [2013]. Fire Sales in a Model of Complexity. *Journal of Finance* 68 (6), 2549-2587.
- Canes-Wrone, B., & Park, J. K. [2014]. Elections, Uncertainty and Irreversible Investment. *British Journal of Political Science*, 44(01), 83-106.
- Caraway, T. L., Rickard, S. J., & Anner, M. S. [2012]. International negotiations and domestic politics: the case of IMF labor market conditionality. *International Organization*, 66[1], 27.

- Carr, M. D., & Jayadev, A. [2015]. Relative income and indebtedness: evidence from panel data. *Review of Income and Wealth*, 61(4), 759-772.
- Carrillo, P. E., Early, D. W., & Olsen, E. O. [2014]. A panel of interarea price indices for all areas in the United States 1982–2012. *Journal of Housing Economics*, 26, 81-93.
- Carroll, C. D. [2003]. Macroeconomic expectations of households and professional forecasters. *the Quarterly Journal of economics*, 269-298.
- Case, K. E., Shiller R.J., Thompson. A. [2012] What have they been thinking? Home buyer behavior in hot and cold markets, *Brookings Papers on Economic Activity* (Fall): 265-298.
- Cheng I., Raina S., Xiong W.[2014].Wall street and the housing bubble. *The American Economic Review* 104.9 (2014): 2797-2829.
- Chernenko, S., Hanson, S. G., & Sunderam, A. [2015]. Who neglects risk? Investor experience and the credit boom. *Unpublished working paper, NBER*, (20777).
- Christen, M., & Morgan, R. M. [2005]. Keeping up with the Joneses: Analyzing the effect of income inequality on consumer borrowing. *Quantitative Marketing and Economics*, 3(2), 145-173.
- Clements, B., Gupta, S., & Nozaki, M. [2013]. What happens to social spending in IMF- supported programmes?. *Applied Economics*, 45[28], 4022-4033.
- Coibion, O., Gorodnichenko, Y., Kudlyak, M., & Mondragon, J. [2014]. Does Greater Inequality Lead to More Household Borrowing? New Evidence from Household Data. *National Bureau of Economic Research*, No. w19850.
- Copelovitch, M. S. [2010]. Master or servant? Common agency and the political economy of IMF lending. *International Studies Quarterly*, 54[1], 49-77.
- Coval, J. D., Jurek J.W., Stafford E., [2009]. The Economics of Structured Finance. *Journal of Economic Perspectives* 23(1): 3–25.
- De Fusco A. [2015] Homeowner Borrowing and Housing Collateral: New Evidence from Expiring Price Controls. *Mimeo, University of Pennsylvania*
- De Giorgi, G., Frederiksen, A., & Pistaferri, L. [2016]. *Consumption Network Effects* . *National Bureau of Economic Research*, No. w22357.
- De Nardi, M., French E., Benson D [2011]. Consumption and the Great Recession. *National Bureau of Economic Research*, No. w17688.

- Devries, P., Guajardo, J., Leigh, D., & Pescatori, A. [2011]. A new action-based dataset of fiscal consolidation. *IMF Working Papers*, 1-90.
- Dreher, A. [2004]. A public choice perspective of IMF and World Bank lending and conditionality. *Public Choice*, 119[3-4], 445-464.
- Dreher, A., & Jensen, N. M. [2007]. Independent actor or agent? An empirical analysis of the impact of US interests on International Monetary Fund conditions. *Journal of Law and Economics*, 50[1], 105-124.
- Dreher, A., Sturm, J. E., & Vreeland, J. R. [2009a]. Development aid and international politics: Does membership on the UN Security Council influence World Bank decisions?. *Journal of Development Economics*, 88[1], 1-18.
- Dreher, A., Sturm, J. E., & Vreeland, J. R. [2009b]. Global horse trading: IMF loans for votes in the United Nations Security Council. *European Economic Review*, 53[7], 742- 757.
- Driscoll, J. C., & Kraay, A. C. [1998]. Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics*, 80(4), 549-560.
- Duesenberry, J.S. [1949]. *Income, Saving and the Theory of Consumer Behavior*. Cambridge, Massachusetts: *Harvard University Press*
- Eeckhout, J., Lindenlaub I., [2015]. Unemployment Cycles. *Institute for Fiscal Studies*, No. W15/26.
- Eggertsson, G. B., & Krugman, P. [2012]. Debt, deleveraging, and the liquidity trap: A Fisher-Minsky-Koo approach. *The Quarterly Journal of Economics*, 127[3], 1469-1513.
- Favara, G., & Imbs, J. [2015]. Credit supply and the price of housing. *The American Economic Review*, 105(3), 958-992.
- FED [2008]. Neither Boom nor Bust: How Houston's Housing Market Differs from the Nation's. January 2008
- Foote C., Gerardi K., Willen P., [2012]. Why Did So Many People Make So Many Ex Post Bad Decisions? The Causes of the Foreclosure Crisis. *Federal Reserve Bank of Boston Public Policy Discussion Paper*, 12-2.
- Frank, R. H., Levine, A. S., & Dijk, O. [2014]. Expenditure Cascades. *Review of Behavioral Economics*, 1(1-2), 55-73.

- Fuster, A., & Vickery, J. [2015]. Securitization and the fixed-rate mortgage. *Review of Financial Studies*, 28(1), 176-211.
- Gennaioli N., Ma Y., Shleifer A. [2015]. Expectations and Investment. *NBER Macroeconomics Annual 2015*, Vol. (30), National Bureau of Economic Research.
- Gennaioli N., Shleifer A., Vishny R., [2015]. Neglected Risk: The Psychology of Financial Crises. *The American Economic Review* vol. 105(5), pages 310-14.
- Gerber A., Huber G. [2009]. Partisanship and Economic Behavior: Do Partisan Differences in Economic Forecasts Predict Real Economic Behavior?. *American Political Science Review* 103 : 407-426.
- Gerber A., Huber G. [2010]. Partisanship, political control, and economic assessments. *American Journal of Political Science* 54.1: 153-173.
- Giannetti, M., & Wang, T. Y. [2014]. Corporate scandals and household stock market participation. *European Corporate Governance Institute (ECGI)-Finance Working Paper*, (405).
- Glaeser, E. L., & Nathanson, C. G. [2015]. An extrapolative model of house price dynamics. National Bureau of Economic Research, WP. No. w21037.
- Glaeser E. [2013]. A nation of gamblers: Real estate speculation and American history. *National Bureau of Economic Research*, WP. No. w18825.
- Gould, E. R. [2003]. Money talks: Supplementary financiers and international monetary fund conditionality. *International Organization*, 57[03], 551-586.
- Gould, E. R. [2006]. Money talks: The International Monetary Fund, conditionality, and supplementary financiers. *Stanford University Press*.
- Greenwood R., Shleifer A. [2014]. Expectations of returns and expected returns. *Review of Financial Studies* 27.3 (2014): 714-746.
- Hoechle, D. (2007). Robust standard errors for panel regressions with cross-sectional dependence. *Stata Journal*, 7(3), 281.
- Iida, K. [1993]. When and how do domestic constraints matter? Two-level games with uncertainty. *Journal of Conflict Resolution*, 37[3], 403-426.
- International Monetary Fund [2003]. *World Economic Outlook*, April 2003.
- Independent Evaluation Office of the International Monetary Fund. [2003]. Fiscal Adjustment in IMF-Supported Programs. .

- Independent Evaluation Office of the International Monetary Fund. [2008]. Structural Conditionality in IMF-Supported Programs.
- Jordà, Ò., Schularick, M., & Taylor, A. M. [2015 a]. Betting the house. *Journal of International Economics*, 96, S2-S18.
- Jordà, Ò., Schularick, M., & Taylor, A. M. [2015 b]. Leveraged bubbles. *Journal of Monetary Economics*, 76, S1-S20.
- Jordà, Ò., Schularick, M., & Taylor, A. M. [2016]. The great mortgaging: housing finance, crises and business cycles. *Economic Policy*, 31(85), 107-152.
- Kahler, M. [1993]. Bargaining with the IMF: Two-level strategies and developing countries. In *Double-edged diplomacy: International bargaining and domestic politics* (No. 25). Univ of California Press. 363- 94.
- Kahn, M. [2011]. Do liberal cities limit new housing development? Evidence from California. *Journal of Urban Economics* 69.2 (2011): 223-228.
- Kahneman D., Tversky A. [1972]. Subjective probability: A judgment of representativeness. *Cognitive Psychology* 3 (3), 430–454.
- Kaplan, G., Mitman, K., & Violante, G. L. [2016]. Non-durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data. *CEPR Discussion Papers* (No. 11255).
- Kentikelenis, A., Karanikolos, M., Papanicolas, I., Basu, S., McKee, M., & Stuckler, D. [2011]. Health effects of financial crisis: omens of a Greek tragedy. *The Lancet*, 378[9801], 1457-1458.
- Kentikelenis, A., Karanikolos, M., Reeves, A., McKee, M., & Stuckler, D. [2014]. Greece's health crisis: from austerity to denialism. *The Lancet*, 383[9918], 748-753.
- Kiyotaki N., Moore J. [1997]. Credit Cycles. *Journal of Political Economy* 105 (2): 211– 248
- Krueger, D., & Perri, F. [2006]. Does income inequality lead to consumption inequality? Evidence and theory. *The Review of Economic Studies*, 73(1), 163-193.
- Kuchler T., Zafar B., [2015]. Personal experiences and expectations about aggregate outcomes. *Federal Reserve Bank of New York Staff Reports* n.748 .
- Kumhof, M., Ranciere, R., Winant, P. [2015]. Inequality, leverage, and crises. *The American Economic Review*, 105(3), 1217-1245.

- Laeven, L., & Valencia, F. [2013]. Systemic banking crises database. *IMF Economic Review*, 61 [2], 225-270.
- Lambertini, L., Mendicino C., Punzi M.T.[2013] Expectation-driven cycles in the housing market: Evidence from survey data. *Journal of Financial Stability* 9.4 (2013): 518-529.
- Ling D.C., Ooi J.T., Le T. [2015]. Explaining house price dynamics: Isolating the role of nonfundamentals. *Journal of Money, Credit and Banking* 47.S1: 87-125.
- Lovell, M.C [1986] Tests of the rational expectations hypothesis. *The American Economic Review* 76.1 (1986): 110-124.
- Maattanen, N., & Tervio, M. [2014]. Income distribution and housing prices: an assignment model approach. *Journal of Economic Theory*, 151, 381-410.
- Madeira, C., & Zafar, B. [2015]. Heterogeneous inflation expectations and learning. *Journal of Money, Credit and Banking*, 47(5), 867-896.
- Malmendier, U., & Nagel, S. [2011]. Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?, *The Quarterly Journal of Economics*, 126(2), 373–416
- Malmendier, U., & Nagel, S. [2016]. Learning from inflation experiences. *The Quarterly Journal of Economics*, 131(1), 53-87.
- Manski, C. F. (2004). Measuring expectations. *Econometrica*, 72(5), 1329-1376.
- Martin, A., & Ventura, J. [2011]. Theoretical notes on bubbles and the current crisis. *IMF Economic Review*, 59(1), 6-40.
- Matlack, J. L., Vigdor, J. L. [2008]. Do rising tides lift all prices? Income inequality and housing affordability. *Journal of Housing Economics*, 17(3), 212-224.
- Mian A., Sufi A., [2009]. The consequences of mortgage credit expansion: evidence from the US mortgage default crisis. *The Quarterly Journal of Economics*, November, 1449- 96
- Mian, A., & Sufi, A. [2010]. The Great Recession: Lessons from Microeconomic Data. *The American Economic Review*, 100(2),
- Mian, A., & Sufi, A. [2011]. House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis. *The American Economic Review*, 101(5), 2132-56.

- Mian, A., Rao K., Sufi A. [2013]. Household balance sheets, consumption, and the economic slump. *The Quarterly Journal of Economics* (128) 1687–1726.
- Mian, A., & Sufi, A. [2014]. What explains the 2007-2009 drop in employment? *Econometrica*, 82(6), 2197-2223.
- Mian, A., Sufi A., Khoshkhou N. [2015] Government Economic Policy, Sentiments, and Consumption. *National Bureau of Economic Research*, No. w21316.
- Mo, J. [1994]. The logic of two-level games with endogenous domestic coalitions. *Journal of Conflict Resolution*, 38[3], 402-422.
- Muth J.F. [1961]. Rational expectations and the theory of price movements. *Econometrica* (1961): 315-335.
- Nooruddin, I., & Simmons, J. W. [2006]. The politics of hard choices: IMF programs and government spending. *International Organization*, 60[04], 1001-1033.
- Osborne, M. J., & Rubinstein, A. [1990]. Bargaining and markets. Academic press.
- Paiella M., Pistaferri L., [2014]. Decomposing the wealth effect on consumption. *Stanford University working paper*.
- Pesaran, M. H., & Weale, M. [2006]. Survey expectations. *Handbook of economic forecasting*, 1, 715-776.
- Piazzesi M., Schneider M., [2009]. Momentum traders in the housing market: survey evidence and a search model. *The American Economic Review Papers and Proceedings* 99(2), 406- 411.
- Piketty, T., & Saez, E. [2014]. Inequality in the long run. *Science*, 344(6186), 838-843.
- Ponticelli, J., & Voth, H. J. (2011). Austerity and anarchy: Budget cuts and social unrest in Europe, 1919-2008. Available at SSRN 1899287.
- Przeworski, A., & Vreeland, J. R. [2000]. The effect of IMF programs on economic growth. *Journal of Development Economics*, 62[2], 385-421.
- Putnam, R. D. [1988]. Diplomacy and domestic politics: the logic of two-level games. *International organization*, 42[03], 427-460.
- Rajan, R. G. [2010], Fault Lines: How Hidden Fractures Still Threaten the World Economy, *Princeton University Press*

- Rickard, S. J., & Caraway, T. L. [2014]. International Negotiations in the Shadow of National Elections. *International Organization*, 68(03), 701-720.
- Saiz, A. [2010]. The geographic determinants of housing supply. *The Quarterly Journal of Economics* 125(3).
- Shiller, R. J. [2007]. *Understanding recent trends in house prices and home ownership*. National Bureau of Economic Research (No. w13553)
- Shiller R., [2015]. Irrational Exuberance. *Princeton University Press*.
- Solon, G., Haider, S. J., & Wooldridge, J. M. [2015]. What are we weighting for?. *Journal of Human Resources*, 50(2), 301-316.
- Soo, C. [2013]. Quantifying animal spirits: news media and sentiment in the housing market. *Ross School of Business Paper 1200*.
- Souleles, N. [2004]. Expectations, heterogeneous forecast errors, and consumption: Micro evidence from the Michigan consumer sentiment surveys. *Journal of Money, Credit and Banking* (2004): 39-72.
- Stone, R. W. [2008]. The scope of IMF conditionality. *International Organization*, 62[04], 589-620.
- Strezhnev, A., & Voeten, E. [2013]. United Nations general assembly voting data. *IQSS Dataverse Network*.
- Tabellini, G., Passarelli F., [2013]. Emotions and political unrest. *CESifo Working Paper Series*, 2013.
- Van Treeck, T. [2014]. Did inequality cause the US financial crisis?. *Journal of Economic Surveys*, 28(3), 421-448.
- Vaubel, R. [1986]. A public choice approach to international organization. *Public Choice*, 51[1], 39-57.
- Vaubel, R., Dreher, A., & Soylu, U. [2007]. Staff growth in international organizations: A principal- agent problem? An empirical analysis. *Public Choice*, 133[3-4], 275-295.
- Vreeland, J. R. [2003]. Why do governments and the IMF enter into agreements? Statistically selected cases. *International Political Science Review*, 24[3], 321-343.
- Wang, C. [2014]. Subjective Home Valuations and the Cross-Section of Housing Returns. Manuscript, *Yale University School of Management*.

Woo, B. [2013]. Disaggregating IMF Conditionality: Comparing Determinants of Fiscal Conditions and Financial Sector Conditions. *Working Paper*.

Woo, J., Bova, M. E., Kinda, M. T., & Zhang, M. Y. S. [2013]. Distributional consequences of fiscal consolidation and the role of fiscal policy: what do the data say? *International Monetary Fund WP*. No. 13-195.

Zhang, C., Jia, S., & Yang, R. [2016]. Housing affordability and housing vacancy in China: The role of income inequality. *Journal of Housing Economics (forthcoming)*.

Zhang, J. [2016]. House price expectations: Unbiasedness and efficiency of forecasters. *Real Estate Economics*, 44(1), 236-257.