Institute for International Economic Studies
Monograph Series No. 111

Stockholm University

Financial Choice and Public Policy

Karin Kinnerud

Doctoral Thesis in Economics at Stockholm University, Sweden 2020

Financial Choice and Public Policy

Karin Kinnerud

Academic dissertation for the Degree of Doctor of Philosophy in Economics at Stockholm University to be publicly defended on Friday 11 September 2020 at 15.00 in Nordenskiöldsalen, Geovetenskapens hus, Svante Arrhenius väg 12.

Abstract

Costly reversals of bad policies: the case of the mortgage interest deduction

This paper measures the welfare effects of removing the mortgage interest deduction under a variety of implementation scenarios. To this end, we build a life-cycle model with heterogeneous households calibrated to the U.S. economy, which features long-term mortgages and costly refinancing. In line with previous research, we find that most households would prefer to be born into an economy without the deductibility. However, when we incorporate transitional dynamics, less than forty percent of households are in favor of a reform and the average welfare effect is negative. This result holds under a number of removal designs.

Monetary policy and the mortgage market

This paper quantifies the role of changes in mortgage interest rates and house prices in the transmission of monetary policy. I build a heterogeneous-agent life-cycle model with housing and long-term mortgage contracts. The illiquid nature of housing gives rise to wealthy hand-to-mouth households, and the existence of mortgage financing allows for households to be both relatively poor and have high exposures to changes in the interest rate. I find that the aggregate response of consumption to a real interest rate shock is highly dependent on the type of mortgage contracts available and the possibility to refinance. In an economy with adjustable-rate mortgages, the consumption response is more than six times as large as compared to when fixed-rate mortgages are used. Hence, a detailed understanding of the contract structures in the mortgage market is an important input into the analysis of monetary policy.

Mortgage lending standards: implications for consumption dynamics

In this paper, we investigate to what extent stricter mortgage lending standards affect households' ability to smooth consumption. Using a heterogeneous-household model with incomplete markets, we find that a permanently lower loan-to-value (LTV) or payment-to-income (PTI) requirement only marginally affects the aggregate consumption response to a negative wealth shock. We show that even the distribution of marginal propensities to consume across households is remarkably insensitive to these permanent policies. In contrast, households' consumption responses can be reduced if a temporary stricter LTV or PTI requirement is implemented prior to a negative wealth shock. However, strong assumptions need to be made for temporary policies to be welfare improving.

Inertia of dominated pension investments: evidence from an information intervention

The market for long-term savings in mutual funds is characterized by high price dispersion between similar funds. In this paper, we conduct an empirical investigation into possible causes of imperfect competition in this market. We discriminate between three main hypotheses on the demand side: a lack of awareness of price dispersion, search costs, and financial illiteracy. We run a large-scale field experiment in the Swedish public pension system, where information letters are sent to savers in two dominated index funds. We show that an information intervention that increases the awareness of a cheaper, dominating fund, at the same time as it reduces the search costs for finding such an alternative, can significantly improve households' real investment allocations. Nonetheless, a majority of savers who are sent information about the name of the dominating fund do not switch funds. Thus, the high degree of inertia in pension investments remains even when search frictions for identifying dominating alternatives are eliminated.

Keywords: *heterogeneous households, housing, mortgage interest deduction, welfare, monetary policy, mortgage contracts, mortgage lending policies, consumption, pensions, field experiment, search costs, inattention, dominated choices, financial literacy.*

Stockholm 2020
http://urn.kb.se/resolve?urn=urn:nbn:se:su:diva-182452

ISBN 978-91-7911-206-6
ISBN 978-91-7911-207-3
ISSN 0346-6892

Department of Economics

Stockholm University, 106 91 Stockholm

Stockholm University

FINANCIAL CHOICE AND PUBLIC POLICY

Karin Kinnerud

Stockholm University

Financial Choice and Public Policy

Karin Kinnerud

ISBN print 978-91-7911-206-6
ISBN PDF 978-91-7911-207-3
ISSN 0346-6892

Cover photo: Picture from rawpixel.com, picture ID 385458
Back-cover photo by Hanna Weitz, Stockholm University
Printed in Sweden by Universitetsservice US-AB, Stockholm 2020

Doctoral Dissertation
Department of Economics
Stockholm University

Abstracts

Costly reversals of bad policies: the case of the mortgage interest deduction

(with Markus Karlman and Kasper Kragh-Sørensen)

This paper measures the welfare effects of removing the mortgage interest deduction under a variety of implementation scenarios. To this end, we build a life-cycle model with heterogeneous households calibrated to the U.S. economy, which features long-term mortgages and costly refinancing. In line with previous research, we find that most households would prefer to be born into an economy without the deductibility. However, when we incorporate transitional dynamics, less than forty percent of households are in favor of a reform and the average welfare effect is negative. This result holds under a number of removal designs.

Monetary policy and the mortgage market

When a central bank changes the interest rate, it affects many households directly through their mortgage interest payments. If these households are constrained in their spending, this channel can have real and direct effects on aggregate demand. However, this channel is absent in standard frameworks of monetary policy. In such frameworks, changes in the policy rate affect consumption demand only via a forward-looking Euler equation. To quantify the mortgage interest rate channel, I build a heterogeneous-agent life-cycle model with housing and long-term mortgage contracts. The illiquid nature of housing gives rise to wealthy hand-to-mouth households, and the existence of mortgage financing allows for households to be both relatively poor and have high exposures to changes in the interest rate. I find that the aggregate response of consumption to a real interest rate shock is highly dependent on the type of mortgage contracts available and the possibility to refinance. In an economy with fixed-payment long-term mortgages, the response of consumption is 50 percent higher due to changes in mortgage interest rates and the endogenous response in house prices. However, in an economy with adjustable-rate mortgages, the consumption response is more than six times as large as compared to when fixed-rate mortgages are used. Hence, a detailed understanding of the contract structures in the mortgage market is an important input into the analysis of monetary policy.

Mortgage lending standards: implications for consumption dynamics
(with Markus Karlman and Kasper Kragh-Sørensen)
In this paper, we investigate to what extent stricter mortgage lending standards affect households' ability to smooth consumption. Using a heterogeneous-household model with incomplete markets, we find that a permanently lower loan-to-value (LTV) or payment-to-income (PTI) requirement only marginally affects the aggregate consumption response to a negative wealth shock. We show that even the distribution of marginal propensities to consume across households is remarkably insensitive to these permanent policies. In contrast, households' consumption responses can be reduced if a temporary stricter LTV or PTI requirement is implemented prior to a negative wealth shock. However, strong assumptions need to be made for temporary policies to be welfare improving.

Inertia of dominated pension investments: evidence from an information intervention
(with Louise Lorentzon)
The market for long-term savings in mutual funds is characterized by high price dispersion between similar funds. In this paper, we conduct an empirical investigation into possible causes of imperfect competition in this market. We discriminate between three main hypotheses on the demand side: a lack of awareness of price dispersion, search costs, and financial illiteracy. We run a large-scale field experiment in the Swedish public pension system. Information letters are sent to pension savers in two index funds, where there exists a cheaper fund with the same index strategy. We show that an information intervention that increases the awareness of a cheaper, dominating fund, at the same time as it reduces the search costs for finding such an alternative, can significantly improve

households' real investment allocations. Nonetheless, a majority of savers who are sent information about the name of the dominating fund do not switch funds. Thus, the high degree of inertia in pension investments remains even when search frictions for identifying dominating alternatives are eliminated.

To my parents

När man sitter så här bland bergen
och pratar med stenar
måste man vara alldeles tyst
Annars svarar de inte

—Okänd

Acknowledgements

This thesis was made possible through the guidance and support of many wonderful people. I am extremely grateful to have had the opportunity to be a part of the Institute for International Economic Studies (IIES). The IIES provides a uniquely embracing environment for PhD students, and I am thankful for having been surrounded by so many great researchers over the past years. It is a true privilege to every day be able to learn, explore, discuss research, and participate in research seminars among excellent thinkers.

I would like to begin by directing a big thank you to my main advisor, Per Krusell. He has provided much more than anyone could ever ask for from an advisor. My way of thinking has been forever influenced by his approach to economics. In the rare instance where he did not immediately have a suggestion or an answer to my problem, he would come back a day or so later with a well thought through advice. Further, he has been supportive and caring, not only about the progress of my research, but also about my general well-being. I am grateful to have had these years to pick his brain, and I hope that he will accept regular calls from me in the future.

Second, it has been a privilege to have Kurt Mitman as a second advisor. A back-of-the-envelope calculation suggests that a 15-minute meeting with him is related to a doubling of my productivity that week. I am particularly thankful that he always quickly provided me with suggested solutions to my problems, and also his generous support during the job market.

My collaborations with other students have been extremely valuable to me. Thank you Markus Karlman, Kasper Kragh-Sørensen, and Louise Lorentzon, for being wonderful co-authors and friends. I have learned a lot from all of you, and you made my PhD so much more enjoyable. I appreciate in particular the economic intuition, the genuine interest, and the positive attitude of Markus; the problem-solving capabilities, the attention to details, and the ambition and hard work of Kasper; and the drive and determination, the creativity, and the enthusiasm about economics of Louise. Most of all I value all the fun times we have had together.

There are a number of researchers that have had a special impact on my development as a PhD student. John Hassler challenges me with questions from other perspectives than I have considered. He is also a role model in his engagement in the policy debate, and I appreciate our talks about the outdoors. Timo Boppart is a great macro economist who I trust to always give me a well thought through perspective on any question I may have. In all of my projects, Roine Vestman has guided me and provided valuable feedback. Jonathan de Quidt, I would like to thank in particular for his commitment and numerous useful suggestions on the paper about pension investments. I would like to thank Arash Nekoei for his encouragements and for always caring. I am also grateful for Mitch Downey's positively delivered constructive feedback and support during the job market. David Strömberg's thoughtful comments and clever suggestions have been extremely valuable. Finally, I would like to thank Lars Calmfors for the interesting discussions about economics and the Swedish policy debate.

Many great researchers contribute to the wonderful atmosphere at the IIES. I owe Ingvild Almås, Tobias Broer, Alexandre Kohlhas, and Kieran Larkin for their valuable feedback, and Tessa Bold, David Schönholzer, and Yimei Zou for their useful advice during the job market. I also wish to thank Konrad Burchardi, Peter Nilsson, Torsten Persson, and Jakob Svensson.

I am thankful to have gotten to know several other excellent researchers throughout these years. In particular, it has been a joy to discuss research with Sergio de Ferra, Rikard Forslid, Paul Klein, Andreas Madestam, Anna Tompsett, and Robert Östling. Kathrin Schlafmann deserves a special thank you for introducing me to the area of quantitative macroeconomics. It is a great benefit to have such an excellent researcher as a teacher. Further, I would like to direct a thank you to the researchers at Sveriges Riksbank, who provided valuable feedback during my internship there and made the visit enjoyable. I am also grateful for Tim Krehbiel's support throughout my master studies.

I wish to express my gratitude to the head of administration at the IIES, Christina Lönnblad, for the excellent support throughout the whole PhD, and in particular while finalizing this thesis. Anne Jensen has also been of great help, and I am thankful for the technical support by Karl

Eriksson. A special thank you goes to Ulrika Gålnander for all the help, and for being such a charming, supportive, and funny friend.

The fourth year of the PhD I spent at Stanford University. I am grateful to Monika Piazzesi and Martin Schneider for all their advice and for including me in their group of students. Their passion for economics is admirable. I am also thankful for the group of macro students at Stanford, in particular Alessandra Peter.

Among my fellow macro students at the IIES I would like to thank Mattias Almgren, Agneta Berge, José-Elías Gallegos Dago, Philipp Hochmuth, Markus Kondziella, John Kramer, Carolina Lindholm, Fredrik Paues, and Markus Peters for many useful and inspiring conversations. It has been a huge benefit to have so many great students around. A special thanks goes to Claire Thürwächter for the many interesting discussions and for being so helpful. Magnus Åhl has been a great support and office mate. Thank you for the numerous discussions and for taking on all my questions about monetary policy. I am also happy to have had Richard Foltyn and Has van Vlokhoven to go to for many of my queries.

In the older cohorts of students, Saman Darougheh, Karl Harmenberg, Matilda Kilström, Hannes Malmberg, Jonna Olsson, Jósef Sigurdsson, and Erik Öberg provide great inspiration, both as excellent researchers and in their passion and curiosity for economics. It was an inclusive and friendly group of macro students to join in the early years of the PhD.

I have had the fortune of several wonderful and bright office mates throughout the PhD, where a great sense of humor has been the common theme. Thank you Divya Dev and others for making the office time so enjoyable.

I am also greatly appreciative of the group of students in my cohort that I have shared many joyful times and lively conversations with, in particular Niklas Blomqvist, Jonas Cederlöf, and Sarah Eichmeyer. Andreas Born and Benedetta Lerva have enriched my life with many laughs and exciting discussions.

Many other fellow PhD students throughout these years have contributed to a stimulating and enjoyable environment. Thanks to all of you, and especially Anna Aevarsdottir, Serena Cocciolo, Montasser Ghachem, Mathias Iwanowsky, Dany Kessel, Roza Khoban, Erik Lindgren, Elisabet Olme, Fabian Sinn, Sreyashi Sen, Tamara Sobolevskaia, and Xueping Sun.

As I made a detour into banking for a few years before starting the PhD, I would like to express my gratitude to Roger Josefsson who, with his inspiring excitement about economics, contributed to my application to the PhD program. I would also like to thank Markus Hveem for his support during the challenging first years of the PhD.

Finally, I would like to thank a number of people who mean a lot to me. Helene Werner, thank you for the many interesting discussions about human behavior and all the laughs throughout the past twenty-something years.

Chris, thank you for giving me energy and for making me laugh when I need it, for joining me to the mountains I dream of, and for enriching my life during this past year.

Lisa, you have always set the bar for me. You started off strong by finishing all the text books for the first three grades already in your first year of school, and ever since then you have shown that the norm is not necessarily what to aim for, and standard is boring. I am also thankful for your honesty and that you always support me, whenever I need it.

My parents have always encouraged me to take on whatever interests me. You have given me the best opportunities in life, and I am thankful for your constant commitment and endless support. I have felt your support regardless of where I have been and what I have been up to. With equal excitement you have learned about the college golf system in the U.S., the Swedish bond market, as well as the economic job market. You have always respected my stubbornness and my priorities, which I am forever grateful for.

Karin Kinnerud
Stockholm, Sweden
August 2020

Contents

Introduction

This thesis consists of four self-contained essays. A common theme across the chapters is the emphasis on how different households are affected differently by public policies. I investigate how including a rich heterogeneity of households in the analysis can influence the main takeaways from policy evaluations as well as our understanding of how the economy as a whole is affected by the policies. The policies considered include a removal of the mortgage interest deductibility, monetary policy, stricter mortgage regulations, and fund choices in pension systems.

Another common theme across the first three papers is the focus on policies in the mortgage and housing markets. The tax treatment and the regulations of the mortgage market often lead to heated debates in many countries. For most households, housing and mortgage choices are the largest and most important financial decisions they make. Furthermore, policies and regulations of the mortgage market often affect households differently. Evaluating such policies therefore requires a rich framework that can capture how different households are affected. Below follows a somewhat less technical summary of the four chapters.

In the first chapter, **Costly reversals of bad policies: the case of the mortgage interest deduction**, jointly written with Markus Karlman and Kasper Kragh-Sørensen, we study how U.S. households are affected by removing the mortgage interest deduction (MID) and whether such a removal is a good idea.

The MID is a tax subsidy that has received a great deal of attention in policy discussions in the U.S. The subsidy allows homeowners to deduct mortgage interest payments from their earnings before paying income taxes. As the MID can reduce the tax payments for homeowners, it effectively lowers the cost of mortgage financing and therefore the cost of owning a house. Thus, many households are affected by the MID, not only in their decision to own as opposed to rent a home, but also when it comes to how large a house to buy. However, the subsidy is often criticized for mainly benefiting high-earners at the expense of other tax payers. Almost half of the deductions go to households in the top 20 percent of the earnings distribution, whereas households in the bottom 20 percent hardly deduct any mortgage interest payments.

To get a better understanding of who would benefit and who would lose from repealing the MID, we perform experiments in a model that is designed to represent the U.S. economy. We begin by analyzing the long-run welfare effects, i.e., we compare if households would prefer to be born into an economy with or without the MID. We find that a vast majority of households would prefer an economy where mortgage interest payments are not deductible. In an economy without the tax subsidy, households with high earnings want smaller houses. This leads to lower prices of owned and rental housing, which is particularly beneficial for low-earning households. Additionally, when the government no longer subsidizes mortgage interest payments, other taxes can be reduced. Whereas only some households benefit from the MID, all households appreciate a lower labor income tax.

Given the large welfare gains of removing the mortgage subsidy in the long run, we proceed by investigating how current households would be affected by a removal. The consequences of a removal are very different for these households. Today, many households have made long-term housing and mortgage decisions based on the premise that they can deduct their mortgage interest payments. When the subsidy is unexpectedly removed, there is a sharp drop in house prices, which hurts the existing homeowners substantially. Further, many households find themselves with too large houses and mortgages, when they can no longer deduct their interest payments. Renters, on the other hand, gain from the reform as they benefit from the fall in house prices.

We find that households are on average worse off by an immediate removal of the MID, and a majority of households are against such a policy. 70 percent of U.S. households own their home and the gains experienced by renters do not exceed the costs among homeowners. Importantly, these results also hold for alternative removal policies where the deductibility is removed gradually or when a removal is preannounced. In fact, under these alternative implementation policies, even fewer households are in favor of a removal. Although more gradual policies alleviate the losses of those hardest hit by the reform, they also make the benefits smaller. Our results thus show that the costs of reverting a bad policy can be substantial — even to the extent that it might not be worthwhile.

In Chapter 2, **Monetary policy and the mortgage market**, I explore the role of changes in mortgage interest rates and house prices for monetary policy. Further, I investigate if the effectiveness of a central bank's policy depends on if households use mortgages with fixed versus adjustable rates.

A principal concern in economics is how a central bank affects the economy by changing the interest rate. In traditional models of monetary policy, a decrease in the interest rate makes households consume more today and save less, since a lower interest rate effectively lowers the price of consumption today as compared to tomorrow. However, there are other ways in which changes in the interest rate affect households. In this paper, I investigate how households respond to a change in the interest rate when this also impacts mortgage interest rates and house prices.

When a central bank changes the interest rate, it affects many households directly through their mortgage interest payments. With lower mortgage interest payments, some households respond by increasing their consumption, which in turn stimulates overall demand in the economy. To what extent households respond to changes in mortgage interest rates depends on if their mortgage payments are affected by the change in the interest rate, and how constrained in their spending households are to begin with. Further, if there is a change in house prices due to the lower interest rate, the wealth of homeowners is affected, which can also influence their consumption choices.

I start by using a model of the U.S. economy, where the most common mortgage contract is the fixed-payment 30-year mortgage. With this type of contract only new mortgages are affected by a change in the mortgage interest rate. Another type of contract, which is common in many countries, is the adjustable-rate mortgage. With this contract type, a change in the mortgage interest rate directly affects the required mortgage payments for existing mortgagors. I continue by studying how consumption responds to a decline in the interest rate when households use this contract type instead.

In the setting where mortgages have fixed rates, I find that when incorporating changes in mortgage interest rates and house prices, there is a larger increase in consumption when a central bank lowers the interest rate. A decline in the interest rate leads to higher house prices, which in combination with lower mortgage interest rates make households increase their consumption. In particular, households who refinance their mortgage increase consumption substantially. Refinancing households tend to be financially constrained, and choose to refinance not only to capture the lower mortgage interest rate but also to take up a larger mortgage. Since house values increase in response to the lower interest rate, households who refinance are able to take up an even larger mortgage, which allows them to increase consumption further.

In an economy with adjustable-rate mortgages, I show that the consumption increase when the central bank lowers the interest rate is over six times larger, as compared to an economy with fixed-rate mortgages. Not only are all mortgagors affected by the lower mortgage interest rate in this setting, but when the interest rate is temporarily reduced, the short-term mortgage interest rate of adjustable-rate contracts decreases more than the long-term rate of fixed-rate mortgages. This further enhances how much households' mortgage payments are affected by the policy. In addition, house prices increase significantly more when the interest rate is reduced, in the setting with adjustable-rate contracts. Once more, households who refinance their mortgage play a key role for the larger consumption responses. With the stronger increase in house prices, refinancing households can take up much larger mortgages, which allows them to increase consumption more.

To summarize, I find that how effectively a central bank can stimulate overall consumption demand is highly dependent on the type of mortgage contracts available and the possibility to refinance. Hence, a detailed understanding of the contract structures in the mortgage market is an important input into the analysis of monetary policy.

In Chapter 3, **Mortgage lending standards: implications for consumption dynamics**, coauthored with Markus Karlman and Kasper Kragh-Sørensen, we investigate whether stricter mortgage lending standards can dampen the fall in consumption during economic downturns. Specifically, we study to what extent mortgage regulations affect how much households change their consumption, when they experience a temporary fall in wealth.

Governments in many countries have implemented stricter mortgage requirements in recent years. These policies are partly motivated by the experiences of the Great Recession, where areas with a higher growth in mortgage debt before the crisis experienced a stronger drop in consumption when the crisis hit. Regulators hope that the new mortgage requirements will make future downturns less severe. However, it is not obvious that the stricter lending standards are successful in stabilizing the economy. One way in which households can avoid a decrease in consumption is exactly by increasing their debt. By restricting the possibility to borrow, households are left with fewer options to cushion a fall in wealth. Therefore, the consumption response may be stronger than without a policy.

In this paper, we use a model to perform experiments where the loan-to-value (LTV) and the payment-to-income (PTI) requirements are made stricter. The LTV limit specifies the maximum mortgage a household can use, as a share of the house value. The PTI constraint limits the size of the mortgage in relation to earnings. In our experiments, we first study a permanent shift of the LTV limit from the current value of 0.90 to 0.70, or the PTI constraint from its current value of 0.28 to 0.18. Then, we explore the same policies, but when they are only implemented temporarily, in a year preceding an economic downturn.

Our first finding is that permanently stricter policies only marginally affect how much households reduce their consumption, when they experience an unexpected fall in wealth. Still, the policies do affect households in important ways. Fewer households own their home, they have less debt, and they save slightly more on average. Crucially, these changes in behavior are such that households' overall ability to handle economic downturns remains virtually unchanged. This result also holds for larger changes in lending standards.

Our second finding is that temporary stricter mortgage standards can successfully reduce the fall in consumption during an economic downturn. A temporary policy prevents some people from buying a house and it makes some households take up smaller mortgages. Therefore, households have more savings available when the economic downturn occurs than they would have had in the absence of the policy. As a result, they end up better prepared to handle the fall in wealth. However, we only find that a temporary policy improves the well-being of households under specific circumstances. First, the economic downturn has to be large. Second, a policymaker needs to have an informational advantage in that she can foresee the downturn, whereas households cannot.

In the fourth chapter, **Inertia of dominated pension investments: evidence from an information intervention**, jointly written with Louise Lorentzon, we examine potential reasons for why pension savers fail to choose funds with the lowest fees among funds with the same investment strategy. Further, we study if information regarding cheaper, comparable options can improve households' investment allocations.

Despite a large set of mutual funds for savers to choose from, there is a lack of competition in the fund market, as characterized by a high price dispersion among comparable funds. Even among funds that follow the same investment strategy, fund companies charge different fees. The fact that savers pay different fees for the same investments contributes to the observed differences in returns to savings across households. Moreover, differences in returns to savings have recently been shown to account for a large fraction of wealth inequality.

In this paper we investigate three potential reasons for why savers prevail with funds that are dominated, i.e., where there is another fund available that is cheaper and has the same investment strategy. First, some savers may be unaware that a cheaper, comparable fund is available. Second, savers may find it too cumbersome to search for a cheaper fund. That savers find it too costly to search for better alternatives is a commonly cited reason for dominated choices in the literature. Third, savers may not fully comprehend the potential gains to be made by switching to a cheaper fund, by underestimating the effect of compound interest.

To study the causes of dominated fund choices, and to explore if information can improve savers' investment allocations, we send information letters to pension savers in two dominated index funds. We find that information that increases the awareness that a cheaper fund is available makes a significant share of savers switch to the cheaper, comparable fund. Information that also tells the savers about the name of the cheaper fund further increases the share who improve their investment allocations. However, clarifying how much more money savers can expect to have in their pension account by the time they retire, if they switch to the cheaper fund, does not lead to an increase in the share of switchers.

Although many pension savers improve their investment allocations when they receive information about a cheaper alternative, a majority of the savers do not change their investments. We conclude that information about comparable funds is useful for many savers, and could be considered when designing the choice architecture of pension systems. However, why so many savers prevail with poor investment choices remains a puzzle, in particular when providing savers with information that removes the costs of finding a better alternative.

Chapter 1

Costly reversals of bad policies: the case of the mortgage interest deduction*

*This paper has been jointly written with Markus Karlman and Kasper Kragh-Sørensen. We are thankful for helpful discussions with Tobias Broer, Jeppe Druedahl, John Hassler, Priit Jeenas, Per Krusell, Virgiliu Midrigan, Kurt Mitman, Monika Piazzesi, Kathrin Schlafmann, Martin Schneider, Roine Vestman, and seminar participants at the 2018 ECB Forum on Central Banking, the 2018 annual congress of the European Economic Association, the 2018 Nordic Summer Symposium in Macroeconomics and Finance, the 2019 ENTER Jamboree at Tilburg University, Norges Bank, the Norwegian Ministry of Finance, Stanford University, Statistics Norway, Stockholm University, the Swedish Financial Supervisory Authority, the Swedish Ministry of Finance, Sveriges Riksbank, and the 2018 Young Economists Symposium at New York University. We gratefully acknowledge funding from Handelsbanken's Research Foundations and Torsten Söderbergs Stiftelse. All errors are our own.

1.1 Introduction

When the mortgage interest deductibility (MID) was passed into law through the Revenue Act of 1913, it was largely insignificant. Hardly any households paid federal income taxes, and those who did predominantly faced a marginal tax rate of only one percent (Ventry, 2010). Today, the MID has become a symbol of the "American dream" of homeownership and reduces the cost of housing for millions of Americans.

The desirability of the MID has recently been called into question. In public discussions, opponents of the MID argue that it is a costly subsidy that does little to help households into the housing market as a disproportionate share of total deductions are claimed by high earners, who would be homeowners regardless (Desmond, 2017).[1] Moreover, the results in the academic literature generally show that most American households would be better off without the MID in the long run.[2]

In this paper, we study how a removal of the MID affects households both in the short and the long run. While our analysis of long-run effects addresses the question whether households would prefer to be born into an economy with or without the MID, the short-run analysis specifically considers the welfare implications of those alive at the time of the removal. The welfare effects may be substantially different in the short run, as current households have already made long-term housing and financing decisions based on the presumption that they can deduct mortgage interest payments.

We find that although the vast majority of households would prefer to be born into a world without the MID, the implementation costs of a removal exceed the benefits. Less than forty percent of current households are in favor of removing the subsidy and the average welfare effect is significantly negative. Interestingly, more gradual removal policies that enable homeowners to adjust their asset holdings before the MID is removed do not increase the support for a removal. These results are robust to including the tax code changes made in the 2017 Tax Cuts and

[1]Total tax expenditures due to the MID are estimated to 63.6 billion dollars in 2017 (JCT, 2017), which is close to the entire annual spending of the Departments of Commerce, Energy, and Justice.

[2]See, e.g., Chambers et al. (2009), Floetotto et al. (2016), Gervais (2002), and Sommer and Sullivan (2018).

Jobs Act. Further, we cannot find a one-time transfer scheme that taxes winners and compensates losers, within the current generation, that leads to a Pareto improvement under any of the policies we consider. Our results thus show that the costs of reverting a bad policy can be substantial — even to the extent that it might not be worthwhile.

To arrive at this conclusion, we study the welfare effects of a removal of the MID through the lens of a life-cycle model with overlapping generations and incomplete markets in which house and rental prices adjust endogenously to clear the housing market. Households can borrow against their house in the form of long-term mortgages. These loans are subject to equity and payment-to-income requirements, and refinancing is costly. The tenure decision is endogenous and there are transaction costs associated with both buying and selling a house. We include the salient features of the U.S. tax code with respect to housing, namely that imputed rents are not taxed and that property taxes and mortgage interest payments are tax deductible. Furthermore, households can choose between itemized deductions and a standard deduction, where the former includes mortgage interest payments. Both deductions are subtracted from earnings that are subject to a progressive tax schedule.

We perform a series of decompositional exercises to better understand: i) why the results in the long run differ so markedly from those in the short run; and ii) why more gradual policies are ineffective in bridging this gap. A natural starting point is to understand why it is beneficial to remove the MID in the long run. We find that the positive welfare results in the long run are due to changes in several equilibrium objects. Households benefit from lower rental and house prices, a lower labor income tax rate, and higher bequests. The direct effect of removing the MID is an increase in the user cost of owning a house for households that itemize deductions. To accommodate the lower housing demand of these households, house and rental prices fall. Reduced prices make rental services more affordable and owned housing more accessible. To ensure tax neutrality, we let the labor income tax be reduced as the government no longer subsidizes mortgage financing. In addition, more bequests are distributed to households as the average net worth goes up. For most households, these positive effects outweigh the direct negative effect of removing the MID.

In our analysis of the transitional dynamics, we begin by studying the effects of an immediate removal and show that the fall in house prices, which increases welfare in the long run, decreases welfare in the short run. Lower house prices reduce housing equity, and thus the wealth of homeowners and the values of bequests. This effect hurts older homeowners in particular. Furthermore, the direct negative effect of increasing the user cost of owner-occupied housing is more prominent, especially for relatively young households that have just entered the housing market and are highly leveraged.

Given that it is beneficial for the lion's share of households to remove the MID in the long run, we explore two alternative policies that are less abrupt and give households time to adjust their asset holdings before the MID is repealed. First, we analyze the effects of linearly reducing the deductible share of mortgage interest payments over fifteen years. Second, we consider an announcement policy in which households can fully deduct their interest payments on mortgages for another fifteen years, after which no payments can be deducted. We find that the immediate policy actually results in the smallest average welfare loss among the policies and has the highest share of households who benefit from a removal. More gradual policies do successfully mitigate the welfare losses of older homeowners and households with large mortgages and high earnings. Importantly, though, these policies also significantly reduce the benefits associated with the immediate policy. Renters prefer reforms in which prices and taxes fall rapidly as they are not directly affected by an MID removal. Higher income and property taxes under more gradual policies also push a considerable share of homeowners that realize welfare gains under an immediate reform into negative welfare territory.

There is a relatively new literature that uses dynamic models with heterogeneous agents to evaluate the consequences of repealing the MID. We build on this strand of the literature, in particular on the work by Floetotto et al. (2016) and Sommer and Sullivan (2018) who both show the importance of studying heterogeneous effects in the implementation phase of housing tax reforms. We contribute to the literature in three ways.

First, contrary to the findings in Floetotto et al. (2016) and Sommer and Sullivan (2018), we find a large and negative average welfare effect of an immediate removal policy and that a majority of households are against such a reform. Although our model shares many similarities with the models in these papers, there are some key differences leading to the discrepancy in the results.[3] Of particular importance is that housing equity is less liquid in our model, due to the refinancing costs of existing mortgages. These costs are considerable, both in the data and in our model, and make it more difficult for households to cushion negative shocks.

Our analysis also differs from that of Sommer and Sullivan (2018) along other important dimensions. We use a model that realistically captures the full life cycle of households and show that the inclusion of retirees is of quantitative importance for the welfare analysis. Specifically, we find that homeowners in retirement are worse off relative to the average working-age household when the MID is removed. For retirees, housing wealth constitutes a greater proportion of total resources, and they have fewer periods left to smooth the negative wealth shock caused by the house price decline. Moreover, in our analysis, households incur negative welfare effects from receiving smaller bequests along the transition due to the sudden house price drop.

Floetotto et al. (2016) study the short-run impact of an MID repeal using a life-cycle model that includes a bequest motive. However, in their analysis, mortgage interest deductions are claimed against earnings that are subject to a proportional labor income tax rate, and all homeowners are implicitly assumed to itemize deductions. In contrast, homeowners in the U.S. and in our model face a progressive labor income tax schedule, and a significant share of households with a mortgage do not itemize deductions. These features allow our model to replicate the pronounced skewness of mortgage interest deduction claims towards high-earning households as seen in the data.

The second contribution of this paper is that we consider and compare the welfare effects of alternative policies for removing the MID. We believe that our analysis of alternative policies enhances the understanding of why the MID has been challenging to repeal, and what type of trade-offs a policymaker faces. Importantly, our results suggest that natural candidates for removal policies – more gradual policies – are not necessarily preferred

[3] In terms of the long-run analysis, we corroborate the important result in Sommer and Sullivan (2018) that homeownership increases when the MID is removed.

by households. Overall, our findings are closely related to those in Conesa and Krueger (1999), who find negative welfare effects of a transition from a pay-as-you-go social security system to a fully funded system, with the highest fraction of households in favor of an immediate reform.

Finally, we contribute by assessing how the 2017 Tax Cuts and Jobs Act affects the welfare consequences of removing the MID. The tax reform substantially reduces the number of households who itemize deductions, as the standard deduction is almost doubled and a cap on deductions for state and local income tax payments and property tax payments is introduced. Although fewer households claim mortgage interest deductions, we find that a majority of households are against a removal and the average welfare effect is still negative in the short run. The MID removal has a more moderate effect on taxes and prices, which reduces the welfare losses for homeowners, but also the welfare gains for renters.

The remainder of the paper is organized as follows. In Section 1.2 we present the model. We explore a simplified version of the model in Section 1.3 and use it to discuss the net benefit of owner-occupied housing and how it is affected by the MID. The calibration of the baseline economy is presented in Section 1.4, along with a comparison to both targeted and non-targeted data moments. Section 1.5 shows and discusses the results of the different policy experiments, while section 1.6 concludes the paper.

1.2 Model

To analyze the effects of removing the mortgage interest deductibility, we construct a life-cycle model with overlapping generations and incomplete markets. The model is in discrete time, where one model period corresponds to three years. It features three types of agents: households, rental firms, and a government. Households start their lives with different levels of net worth. Further heterogeneity arises from aging and idiosyncratic earnings shocks. Rental firms operate in a competitive market with free entry and exit, and provide rental services to households. The government taxes households and rental firms in a manner that mimics the U.S. tax system. Importantly, we include the main features of the U.S. tax code with respect to housing, namely that imputed rents are not taxed, and that property taxes and mortgage interest payments are tax deductible.

Furthermore, itemized and standard tax deductions are available to households, and are deducted from earnings that are subject to a progressive tax schedule.

There are three assets in the economy: houses, mortgages, and risk-free bonds. Houses are available in discrete sizes, and there are transaction costs associated with both buying and selling a house. The stock of housing is fixed in aggregate, but flexible in its composition.[4] In equilibrium, house prices and rental prices adjust to clear the housing market. The interest rates on mortgages and bonds are exogenous and the supply of both assets is perfectly elastic.

1.2.1 Households

Households are born with initial assets as in Kaplan and Violante (2014). Over the course of the life cycle, households are hit by idiosyncratic permanent and transitory earnings shocks. A household retires with certainty after period J_{ret} and cannot live past period J. The probability of surviving between any two ages j and $j+1$ is $\phi_j \in [0,1]$, and the agents discount exponentially with a factor β. In each period, a household derives utility from a consumption good c and housing services s through a CRRA utility function with a Cobb-Douglas aggregator

$$U_j(c,s) = e_j \frac{\left(c^{\alpha} s^{1-\alpha}\right)^{1-\sigma}}{1-\sigma}, \tag{1.1}$$

where e_j is an age-dependent utility shifter that captures changes in household size over the life cycle (see, e.g., Kaplan et al. (2020)). There is also a warm-glow bequest motive similar to De Nardi (2004), given by the bequest function

$$U^B(q') = \upsilon \frac{(q' + \bar{q})^{1-\sigma}}{1-\sigma}, \tag{1.2}$$

where υ is the weight assigned to the utility from bequests, q' is the net worth of the household, and $\bar{q}$ captures the extent to which bequests are luxury goods. The objective of the household is to maximize the expected

[4]The main focus of this paper is the short-run effects of a housing subsidy removal. Therefore, we find the assumption of a fixed aggregate supply of housing reasonable.

sum of discounted lifetime utility.

A household enters each period j with bonds b, mortgage m, and house h, according to the choices made in the previous period. In the current period, earnings y are realized, the household receives bequests, and pays taxes Γ. It then chooses consumption c, housing service s, bonds b', mortgage m', and house h'. Housing services are either obtained via the agent's owned house or from a rental company. Each unit of housing costs p_h to buy and p_r to rent. An owned house of size h' produces housing services through a linear technology $s = h'$. These services have to be consumed by the owner of the house, which implies that households cannot be landlords. We model landlords implicitly through a rental market, as landlords are treated as business entities in the U.S. tax code. In addition, since landlords are treated as businesses, they are not directly affected by a removal of the mortgage interest deductibility. Households can use mortgages m', with the interest rate r^m, to finance their homeownership. Bonds b' can be purchased in any non-negative amount, earning interest $r < r^m$.

Mortgages are long-term and non-defaultable. In each period, a homeowner with a mortgage needs to adhere to an amortization schedule that specifies a minimum payment $\chi_j m$, where χ_j is defined as

$$\chi_j = \left(\sum_{k=1}^{M_j} \left[\frac{1}{(1+r_m)^k} \right] \right)^{-1}. \tag{1.3}$$

The maturity of the mortgage is given by $M_j = \min\{10, J-j\}$, which implies that the minimum payment is similar to that of an annuity mortgage with either 30 years remaining (10 model periods) or the number of years until the households dies with certainty.[5] A household that stays in a given house has the option to not follow the repayment plan by taking up a new mortgage, but then it incurs a fixed refinancing cost ς^r.

A household that takes up a new mortgage, either when it purchases a new house or refinances an existing mortgage, has to comply with two constraints. First, a loan-to-value (LTV) requirement states that a household can only use a mortgage to finance up to an exogenous share

[5] The 30-year mortgage contract is the most common plan in the U.S. For other ways of modeling long-term mortgages, see, e.g., Kaplan et al. (2020) or Boar et al. (2020).

$1-\theta$ of the house value

$$m' \leq (1-\theta)p_h h'. \tag{1.4}$$

Second, a payment-to-income (PTI) constraint ensures that a household can only choose a mortgage such that the cost of housing-related payments does not exceed a fraction ψ of current permanent income z. Formally,

$$\chi_{j+1} m' + (\tau^h + \varsigma^I)p_h h' \leq \psi z, \tag{1.5}$$

where τ^h and ς^I capture property tax and home insurance payments, respectively.[6] The PTI and LTV requirements together with the refinancing cost limit the possibility to extract housing equity. Thus, instead of paying off a mortgage to increase the housing equity, liquid bonds constitute a more suitable instrument for precautionary savings purposes. In equilibrium, some households will therefore choose to hold bonds and mortgages at the same time.

The household problem has five state variables: age j, permanent earnings z, mortgage m, house size h, and cash-on-hand x. The first two are exogenous, while the latter three are affected by a household's choices. State x is defined as

$$x \equiv y + (1+r)b - (1+r^m)m + (1-\varsigma^s)p_h h - \delta^h h + a - \Gamma, \tag{1.6}$$

where $(1-\varsigma^s)p_h h$ is the value of the house net of transaction costs.[7] The transaction cost of selling a house is modeled as a share ς^s of the house value. The maintenance cost $\delta^h h$ is paid by all homeowners, and is proportional to the size of the house. Initial assets and inheritance are captured by the term a. For a detailed description of how inheritance is modeled, see Section 1.2.3. Total tax payments are represented by Γ, and

[6] Mortgage payments, property taxes, and home insurance costs are three main components used by banks to assess the payment capability of mortgage applicants. The home insurance payment does not enter the household budget constraint in the model, but is included in the PTI requirement for calibration purposes, see Section 1.4.1.

[7] For computational reasons, and without loss of generality, we define cash-on-hand as including the net revenue of selling the house. Households who do not sell their house between any two periods do not incur any transaction costs.

consist of five different taxes

$$\Gamma \equiv \tau^l y + I^w \tau^{ss} y + \tau^c rb + \tau^h p_h h + T(\tilde{y}). \tag{1.7}$$

Similar to the U.S. tax system, a household pays a local labor income tax τ^l, a payroll tax τ^{ss} (only paid by working-age households, represented by the dummy variable I^w), a capital income tax τ^c, a property tax on owned housing τ^h, and a federal labor income tax $T(\tilde{y})$.[8] The federal labor income tax is given by a non-linear tax and transfer system, which is a function of earnings net of deductions $\tilde{y}$. In turn, deductions depend on a household's mortgage, house value, and gross earnings. For a detailed description of the non-linear tax and transfer system see section 1.2.3, in particular equations (1.10) and (1.11).

The household problem includes the discrete choice of whether to rent a home, buy a house, stay in an existing house but refinance the mortgage, or stay in an existing house and follow the repayment plan. Therefore, we split the household problem into these four respective cases, and solve it recursively. Let us define the expected continuation value in the next period as

$$\mathbb{E}\left[W_j(z', x', h', m', q')\right] \equiv \phi_j \mathbb{E}\left[V_{j+1}(z', x', h', m')\right] + (1 - \phi_j) U^B(q').$$

If the household chooses to rent, the optimization problem is given by

$$V_j^R(z, x) = \max_{c,s,b'} U_j(c, s) + \beta \mathbb{E}\left[W_j(z', x', h', m', q')\right]$$

subject to

$$\begin{aligned}
x' &= y' + (1 + r)b' + a' - \Gamma' \\
q' &= b' \\
x &= c + p_r s + b' \\
s &\in S \\
c &> 0, h' = 0, b' \geq 0, m' = 0.
\end{aligned}$$

The problem is characterized by the Bellman equation, the law of motion

[8] The local labor income tax is mainly included to ensure that high-earning households are more prone to itemize deductions.

for cash-on-hand, the equation for bequests, the budget constraint where the current period cash-on-hand is given, and a number of additional constraints. In this first case, the household rents a house and can therefore not take up a mortgage, implying $h' = m' = 0$. The choice of housing service is restricted to the ordered set of discrete sizes $S = \{\underline{s}, s_2, s_3, ..., \bar{s}\}$.

If the household chooses to buy a house of a different size than what it entered the period with, such that $h' \neq h$, the problem becomes

$$V_j^B(z, x) = \max_{c,h',m',b'} U_j(c, s) + \beta \mathbb{E}\left[W_j(z', x', h', m', q')\right]$$

subject to

$$\begin{aligned}
x' &= y' + (1 + r)b' + a' - \Gamma' - (1 + r^m)m' + (1 - \varsigma^s)p_h' h' - \delta^h h' \\
q' &= b' + p_h h' - m' \\
x &= c + (1 + \varsigma^b)p_h h' + b' - m' \\
h' &\in H \\
c &> 0, s = h', b' \geq 0, m' \geq 0,
\end{aligned}$$

along with the LTV constraint (1.4), and the PTI constraint (1.5). Since the household in this case buys a house, the budget constraint allows for the use of a mortgage to finance expenditures. The parameter ς^b captures the transaction cost of buying a house, which is modeled as proportional to the house value. Moreover, the household's choice of housing is limited to a set H, which is a proper subset of S. Specifically, the smallest house size $\underline{h}$ in H is larger than the smallest available size in S.[9] Above and including that lower bound, both sets are identical.

If the household decides to stay in the same house as when entering the period, such that $h' = h$, but chooses to refinance its mortgage, the problem is given by

$$V_j^{RF}(z, x, h) = \max_{c,m',b'} U_j(c, s) + \beta \mathbb{E}\left[W_j(z', x', h', m', q')\right]$$

[9] A minimum size of owner-occupied housing $\underline{h}$ is also assumed in, e.g., Cho and Francis (2011), Floetotto et al. (2016), Gervais (2002), and Sommer and Sullivan (2018).

subject to

$$\begin{aligned}
x' &= y' + (1+r)b' + a' - \Gamma' - (1+r^m)m' + (1-\varsigma^s)p_h'h' - \delta^h h' \\
q' &= b' + p_h h' - m' \\
x &= c + b' + (1-\varsigma^s)p_h h - m' + \varsigma^r \\
c &> 0, s = h' = h, b' \geq 0, m' \geq 0,
\end{aligned}$$

along with the LTV constraint (1.4), and the PTI constraint (1.5). In this case, the house size h enters as a state variable in the Bellman equation, since it directly determines the housing choice h'. Moreover, since x is defined such that it includes the value of the house when sold, the budget constraint is corrected for the agent not selling the house. This is done by adding $(1-\varsigma^s)p_h h$ to the expenditures in the budget constraint. The refinancing cost is captured by ς^r.

Finally, if the household decides to stay in its house and follow the repayment plan, the problem is

$$V_j^S(z,x,h,m) = \max_{c,m',b'} U_j(c,s) + \beta\mathbb{E}\left[W_j(z',x',h',m',q')\right]$$

subject to

$$\begin{aligned}
x' &= y' + (1+r)b' + a' - \Gamma' - (1+r^m)m' + (1-\varsigma^s)p_h'h' - \delta^h h' \\
q' &= b' + p_h h' - m' \\
x &= c + b' + (1-\varsigma^s)p_h h - m' \\
m' &\leq (1+r_m)m - \chi_j m \\
c &> 0, s = h' = h, b' \geq 0, m' \geq 0.
\end{aligned}$$

The mortgage level m now enters as an additional state variable as it determines the choice set for m'. Importantly, by following the repayment plan, the household is not subject to the LTV and PTI requirements.

The solution to the household problem is provided by

$$V_j(z,x,h,m) = \max\left\{V_j^R(z,x), V_j^B(z,x), V_j^{RF}(z,x,h), V_j^S(z,x,h,m)\right\}, \quad (1.8)$$

with the corresponding set of policy functions

$$\Big\{c_j(z,x,h,m), s_j(z,x,h,m), h'_j(z,x,h,m), m'_j(z,x,h,m), b'_j(z,x,h,m)\Big\}.$$

1.2.2 Rental market

The rental price p_r is determined in a competitive rental market. This market consists of a unit mass of homogeneous rental firms. Each firm f chooses either to buy a stock of housing h_f at price p_h per unit and rent it out to households, or to invest the value $p_h h_f$ in risk-free bonds. The present value of after-tax profits in the former case is

$$\pi_f^{Rent} = (1-\tau^c)\Big(p_r h_f - \frac{1}{1+\tilde{r}}\Big[\delta^r + \tau^h p'_h + \Delta p'_h\Big] h_f\Big).$$

Firm f's revenue is given by its rental income $p_r h_f$. The firm can deduct its operating expenses from these revenues before paying taxes at the rate τ^c. The operating expenses comprise a maintenance cost $\delta^r > \delta^h$ per unit of housing, a property tax on the value of the rental stock in the next period $\tau^h p'_h h_f$, and any negative price return on the rental stock $\Delta p'_h h_f$, where $\Delta p'_h \equiv p_h - p'_h$.[10] All operating expenses are discounted, as these costs are realized in the next period, at a rate given by the after-tax return on bonds $\tilde{r} \equiv (1-\tau^c) r$.

In case firm f instead invests in bonds, the present value of after-tax profits is given by

$$\pi_f^{Bonds} = \frac{(1-\tau^c)}{1+\tilde{r}} r p_h h_f.$$

Imposing a free entry and exit condition, such that $\pi_f^{Rent} = \pi_f^{Bonds}\ \forall f$, the equilibrium rental price is

$$p_r = \frac{1}{1+\tilde{r}}\Big[\delta^r + r p_h + \tau^h p'_h + \Delta p'_h\Big]. \tag{1.9}$$

[10]The assumption that $\delta^r > \delta^h$ is one common way in the literature to incorporate an advantage of owning (see, e.g., Piazzesi and Schneider (2016)). It was first introduced in Henderson and Ioannides (1983), and can be thought of as representing a moral hazard problem between owners of rental units and their tenants. An alternative approach would be to assume that owned housing units provide more housing services than rental units.

Admittedly, the rental market can be modeled in other ways. This formulation captures that the return of rental investments should be closely related to the return of other assets. An additional advantage of using this approach is that it yields a tractable closed-form solution for the rental price and the net benefit of owning (see equation (1.16)), which is key to understanding how the MID affects the demand for owner-occupied housing.

1.2.3 Government

The role of the government in the model is to provide retirement benefits to households, collect bequests and distribute these to surviving households, and tax the agents in a manner that replicates the U.S. tax system. Households pay five different taxes. The local level labor income tax, the payroll tax, the capital income tax, and the property tax are modeled linearly, as shown in equation (1.7). In contrast, the federal labor income tax is given by a function that mimics the U.S. federal tax and transfer system. The labor income tax function takes earnings net of deductions $\tilde{y}$ as its argument and is assumed to be continuous and convex, following Heathcote et al. (2017). Specifically,

$$T(\tilde{y}) = \tilde{y} - \lambda \tilde{y}^{1-\tau^p}, \tag{1.10}$$

where λ governs the tax level, and τ^p determines the degree of progressivity.

The type and amounts of deductions a household takes affect taxable earnings. Before retirement, households can itemize deductions, opt for the standard deduction, or not deduct at all. Itemized deductions, including mortgage interest payments, are only permissible as long as the sum of these exceeds the standard deduction. During retirement, households can only use the standard deduction or not deduct at all. To summarize, households' taxable earnings are such that $T(\tilde{y})$ is minimized, subject to

$$\tilde{y} \in \begin{cases} \{\max(y - ID, 0), \max(y - SD, 0), y\}, & \text{if } j \leq J_{ret} \text{ and } ID > SD \\ \{\max(y - SD, 0), y\}, & \text{otherwise} \end{cases} \tag{1.11}$$

$$\text{where } ID = \tau^m r^m m + \tau^h p_h h + \tau^l y.$$

The max operators reflect the fact that taxable earnings must be non-negative. SD is the common exogenous amount that can be deducted if households opt for the standard deduction, while ID is the sum of itemized deductions that includes mortgage interest payments, property tax payments, and local tax payments. These are among the most important deductions in the U.S. tax code (Lowry, 2014). The parameter τ^m is the mortgage deductibility rate in the economy and it is the parameter of interest in this paper. In line with the U.S. tax code, τ^m is set to one in the benchmark model. In other words, all mortgage interest payments are deductible from earnings when calculating taxable earnings for an itemizing household. From equations (1.6), (1.7), (1.10), and (1.11), we see that the MID reduces taxable earnings, and hence increases cash-on-hand, provided that the agent itemizes tax deductions and has a mortgage.

Rental firms pay two taxes: the property tax on their rental stock and the capital income tax on their accounting profits. In total, the government's tax revenues from households and rental firms are given by

$$TR = \sum_{j=1}^{J} \Pi_j \int_0^1 \Gamma_{ij} \, di + \int_0^1 \left(\tau^c r h_f + \tau^h p_h h_f\right) \, df, \tag{1.12}$$

where i indexes households, f indexes rental firms, Π_j is the age distribution of households, and Γ are total taxes as defined in equation (1.7). We assume that both households and rental firms are of unit measure. The government uses part of the tax revenues to finance the retirement benefits. The remaining revenues are allocated to spending that does not affect the other agents.

The government collects bequests in the form of bonds, houses, and mortgages from households who die. After the government has received these bequests, it earns the interest on bond holdings, sells the houses and incurs the transaction costs of selling, and pays off any outstanding mortgages including interest. Thus, the net amount collected from households is given by

$$BQ = \sum_{j=1}^{J} \Pi_j (1-\phi_j) \int_0^1 \left((1+r)b'_{ij} + (1-\varsigma^s)p'_h h'_{ij} - (1+r_m)m'_{ij}\right) di. \tag{1.13}$$

In the initial economy with MID, the government distributes some of these bequests to cover the initial asset holdings of newborns, whereas the remainder is, for simplicity, assumed to cover wasteful government spending. Thus, in the initial steady state, inheritance a in equation (1.6) is zero for all households of age $j > 1$.

Altering the MID is likely to affect the amount of bequests left behind. To capture the welfare effects of changes in the bequests collected, we assume that any increase or decrease in bequests is distributed to surviving households (except newborns) in proportion to a household's permanent earnings in the previous period, i.e., $a_j = \gamma z_{j-1}$ for $j > 1$. Specifically, the parameter γ is adjusted such that the amount distributed equals the change in bequests collected.

1.2.4 Equilibrium

In the equilibrium of the model, house and rental prices are endogenously determined and they adjust to ensure that the demand for housing equals the supply of housing. The model setting can be interpreted as a small open economy, where houses can only be purchased by residents and the interest rates on risk-free bonds and mortgages are taken as given.

In the initial steady state with MID, i.e., $\tau_m = 1$, we set the house price p_h equal to one. House values (price times size) are comparable to the data as the supply of housing quantity (size) is perfectly elastic and households' preferences ensure that a realistic share of expenditures is spent on housing. With the house price at hand, the rental price p_r is easily computed from equation (1.9). The rental market clears automatically as we let the rental companies cater any demand for rental units. Taking house and rental prices as given, we solve for the value and policy functions of the households and proceed by simulating the economy. The aggregate housing supply is then given by the overall demand for housing services. In the remainder of the analysis, the housing supply is fixed at this initial level, but its composition is flexible.

When we solve for the steady-state equilibrium without MID, i.e., $\tau_m = 0$, the demand for housing is affected and the house and rental prices adjust to clear the housing market. Further, we solve for the average labor income tax rate λ, such that the government's tax revenues are the same as in the initial steady state, and the bequest rate γ, such that any changes in bequests left behind are distributed to the households. Additionally, a change in the house price affects the purchasing power of a household that receives bequests. To capture the change in purchasing power, the net worth q' that enters the utility function for bequests is deflated by a price index $\alpha + (1-\alpha)p_h$.

To compute a transitional equilibrium, we first choose a sequence of mortgage interest deductibility parameters $\{\tau_t^m\}_{t=1}^{t=T}$, where T is the last transition period. We then solve for the sequences of house and rental prices, $\{p_{ht}, p_{rt}\}_{t=1}^{t=T}$, and the sequences of the parameters governing the average labor income tax rate $\{\lambda_t\}_{t=1}^{t=T}$ and the bequest rate $\{\gamma_t\}_{t=1}^{t=T}$, such that for all $t \in \{1, ..., T\}$, total housing demand equals the initial housing stock, tax neutrality is achieved, and any changes in bequests are distributed to the households. In the transition, the removal policies are implemented unexpectedly and households have perfect foresight of the transition paths of the deductibility parameter, house and rental prices, as well as the tax and bequest parameters. Any unexpected change in the house price in the first period of the transition, affects the profits of the rental companies. We assume that any profit changes in the first period of the transition are distributed to the homeowners in proportion to their cash-on-hand x. For a detailed description of the equilibrium definitions, the computational methods, and the solution algorithms, see the Appendices.

1.3 The MID and the benefit of owning

To better grasp the mechanisms behind the results in this paper, it is useful to understand why households want to own a house in the model and how this is affected by the MID. Our discussion builds upon previous work on the user cost of owning by, e.g., Díaz and Luengo-Prado (2008), but here we distinguish between those who itemize deductions and those who do not, as this is central to our analysis. We compare a household who owns a house of size h' to a similar household who instead obtains the equivalent housing service $s = h'$ on the rental market. The ex-post net benefit of owning NB^{Own}, in any period, is given by

$$NB^{Own} = UC^{Rent} - UC^{Own}, \tag{1.14}$$

where UC^{Rent} is the user cost of renting and UC^{Own} is the user cost of owning. Intuitively, the net benefit of owning is positive whenever owning is less costly as compared to renting.

The user cost of renting is given by $p_r s$, i.e., the rental price times the size of the rental unit. The user cost of owning is more complicated, as an owned house is an asset that comes with the possibility of debt financing. To keep the analysis in this section tractable, we make a few simplifying assumptions as compared to the full model. First, we abstract from any risk by assuming that prices are constant over time and that the earnings in the next period y' are known. Second, we assume that the interest rate on mortgages r^m is equal to the risk-free rate r. Third, we abstract from the possibility of selling and buying a house and hence, from the transaction costs that occur when doing so. Fourth, we assume that local labor income taxes are not tax deductible.

Given the modifications to the full model, the user cost of owning includes the sum of four costs. First, there is a maintenance cost of $\delta^h h'$. Second, there is an opportunity cost of equity. If the equity had not been invested in the house, it would have yielded an after tax return of $\tilde{r}(p_h h' - m')$, where $\tilde{r} \equiv (1-\tau^c)r$ is the net of tax risk-free rate. Third, a homeowner needs to pay a property tax on the house. This property tax cost is modeled as a fixed share of the house value, and is given by $\tau^h p_h' h'$. Last, a homeowner incurs a cost whenever it uses a mortgage to finance its dwelling. The borrowing cost is simply the interest payment on the mortgage rm'.

The costs of owner-occupied housing can be reduced whenever a homeowner chooses to itemize deductions rather than simply opt for a standard deduction. The sum of the itemized deductions amounts to $ID' = \tau^h p_h' h' + \tau^m r m'$, and is subtracted from earnings which, in turn, are subject to the progressive tax schedule $T(\tilde{y}')$. Importantly, any itemized deductions in excess of the standard deduction reduce the tax liabilities of the homeowner and therefore lower the effective cost of property taxes and mortgage financing. The total benefit from being able to itemize deductions is given by

$$I^d \int_{SD}^{ID'} T_{\tilde{y}'}(y' - \hat{D}) d\hat{D},$$

where I^d is an indicator variable for itemized tax deductions. The user cost of owning is the present value of the sum of all costs, adjusted for deductions

$$UC^{Own} = \frac{1}{1+\tilde{r}} \left(\delta^h h' + \tilde{r}(p_h h' - m') + \tau^h p_h' h' + rm' - I^d \int_{SD}^{ID'} T_{\tilde{y}'}(y' - \hat{D}) d\hat{D} \right). \quad (1.15)$$

Substituting equations (1.9) and (1.15) into (1.14), we get

$$NB^{Own} = \frac{1}{1+\tilde{r}} \left[(\delta^r - \delta^h) h' + \tau^c r(p_h h' - m') + I^d \int_{SD}^{ID'} T_{\tilde{y}'}(y' - \hat{D}) d\hat{D} \right]. \quad (1.16)$$

The first term is the benefit of owning due to a lower depreciation of owned housing as compared to rental housing. The second term is the benefit of investing equity in an asset (housing) where the return is not taxed, compared to investing in bonds where the return is taxed at a rate τ^c. This benefit to owner-occupied housing arises because the imputed rent is not taxed. The last term consists of the tax benefits of owner-occupied housing due to property tax and mortgage interest deductions. Thus, the above measure of the net benefit of owning encapsulates the main features of the U.S. tax treatment of housing.

To see how the net benefit of owning is affected by the deductibility parameter τ^m, it is useful to take the derivative of equation (1.16) with respect to mortgages

$$NB_{m'}^{Own} = \frac{1}{1+\tilde{r}} \left[-\tau^c r + I^d T_{\tilde{y}'}(y' - ID') \tau^m r \right]. \quad (1.17)$$

An increase in the mortgage level, and consequently a reduction in equity, has two effects on the net benefit. On the one hand, the reduction in equity means a smaller benefit resulting from the lack of taxation of imputed rent, which is captured by the first term. On the other hand, since mortgage interest payments are tax deductible ($\tau^m = 1$ in the initial steady state), the increased mortgage results in larger deductions and hence a higher net benefit.

Overall, equations (1.16) and (1.17) are key to understanding how the MID affects the net benefit of owning and, subsequently, the demand for owner-occupied housing. First, the MID increases the net benefit of owning by decreasing the cost of mortgage financing only for those who itemize deductions. In the full model, itemizing households are those with relatively large mortgages, houses, or earnings, or a combination of the three. Second, the net benefit of owning due to mortgage interest deductions is increasing in the marginal tax rate. Figure 1.2 illustrates that the marginal tax rate differs substantially between households, leading to significant differences in the user cost of owning between households. Third, the net benefit of owning is positive regardless of the MID, due to the difference in the depreciation rates, the lack of taxation of the imputed rent, and the property tax deduction. In the full model, transaction costs, borrowing constraints, the mortgage interest spread, and the minimum size of owner-occupied housing hinder some households from owning and make some households prefer renting.

1.4 Calibration

We calibrate the model to the U.S. economy. To avoid capturing business-cycle movements in the data, calibration figures are taken from pooled data over the period 1989 - 2013, subject to data availability. Most of our parameters are calibrated independently, based on data or previous studies, whereas the remaining parameters are calibrated using simulated method of moments.

1.4.1 Independently calibrated parameters

Yearly parameter values taken from other studies or calculated directly from the data are listed in Table 1.1.

Demographics and preferences

The households enter the economy at age 23. The probability of a household dying between two consecutive ages is taken from the Life Tables for the U.S. social security area 1900-2100 (see Bell and Miller (2005)). We use the observed and projected mortality rates for males born in 1950. In the model, the retirement age is set to 65, and we assume that all households are dead by the age of 83. Using data from the Panel Study of Income Dynamics (PSID), we specify the equivalence scale e_j as the square root of the predicted values from a regression of family size on a third-order polynomial of age. In the CRRA utility function, we set

Parameter	Description	Value
σ	Coefficient of relative risk aversion	2
τ^l	Local labor income tax	0.05
τ^c	Capital income tax	0.15
τ^{ss}	Payroll tax	0.153
τ^h	Property tax	0.01
τ^m	Mortgage interest deductibility	1
r	Interest rate	0.03
κ	Yearly spread, mortgages	0.014
γ	Bequest rate	0
θ	Down-payment requirement	0.20
ψ	Payment-to-income requirement	0.28
δ^h	Depreciation, owner-occupied housing	0.03
ς^I	Home insurance	0.005
ς^b	Transaction cost if buying house	0.025
ς^s	Transaction cost if selling house	0.07
ς^r	Refinancing cost	3.0
R	Replacement rate for retirees	0.5
B^{max}	Maximum benefit during retirement	51.1

Table 1.1: Independently calibrated parameters, based on data and other studies

Note: The table presents calibrated parameter values. The values are annual for relevant parameters. When simulating the model, we adjust these values to their three-year (one model period) counterparts. The refinancing cost ς^r and the maximum benefit during retirement B^{max} are in 1000's of 2013 dollars.

the coefficient of relative risk aversion σ to 2, which is widely used in the literature.

Assets and bequests

The initial asset holdings for households are calibrated as in Kaplan and Violante (2014). We divide households aged 23-25 in the Survey of Consumer Finances (SCF) into 21 groups based on their earnings. For each of these groups, we calculate the share with asset holdings above 1,000 in 2013 dollars and the median asset holdings conditional on having assets above this limit. The median asset value for each group is scaled by the median earnings among working-age households (23-64) in the SCF data. For model purposes, we rescale these asset values with the median earnings of working-age households in our model.

The parameter γ, which determines how much bequests each household receives, is set to zero in the initial steady state. When conducting the policy experiments, this parameter is adjusted to account for changes in bequests.

Tax system

The local labor income tax rate τ^l is set to 0.05, which is the average state and local labor income tax rate for itemizers in 2011 (Lowry, 2014). The capital income tax τ^c is set to 0.15, to match the maximum rate that applies to long-term capital income for most taxpayers. In the U.S., the payroll tax is levied equally on both the employer and the employee, and amounts to 15.3 percent of earnings (Harris, 2005). Since there is no explicit production sector in our model, we let the full tax burden fall on the worker by setting τ^{ss} to 0.153. The American Housing Survey (AHS) shows that the median amount of real estate taxes per $\$1,000$ of housing value is approximately 10 dollars.[11] Following this estimate, we set the property tax parameter τ^h to 0.01.

The mortgage interest deductibility rate τ^m is our parameter of interest. In the analysis we alter this parameter from one to zero, where the benchmark economy is characterized by full deductibility (τ^m=1).

Market setting

The interest rate is estimated from market yields on the 30-year constant maturity nominal Treasury securities, deflated by the year-to-year headline Consumer Price Index (CPI). The average real rate over the period 1997 to 2013 is 3.4 percent (Federal Reserve Statistics Release, H15, and the Bureau of Labor Statistics, Databases & Tables, Inflation & Prices). We set the real interest rate r to 0.03. Using the Federal Reserve's series of the contract rate on 30-year fixed-rate conventional home mortgage commitments over the period 1997 to 2013, we find that the average yearly spread to the above Treasuries is 1.4 percentage points. Consequently, we choose a yearly spread for mortgages κ of 0.014, implying a mortgage interest rate r^m of 0.044.

Similar to Floetotto et al. (2016) and Sommer and Sullivan (2018), we set the minimum down-payment requirement θ to 0.20 in the model. The

[11] See table C-10-OO in the 2011 and 2013 American Housing Survey, and table 3-13 in the 2009 wave.

payment-to-income requirement ψ is taken from Greenwald (2018) and is set to 0.28.

The depreciation rate of owned housing is set to 3 percent. This follows from the estimate of the median depreciation rate of owned housing, gross of maintenance, in Harding et al. (2007). The transaction costs of buying and selling a house are taken from Gruber and Martin (2003). They use the median transaction costs from CES data and estimate the costs of buying and selling to be 2.5 and 7 percent of the house value, respectively. The home insurance is calibrated to match the median property insurance payment in the AHS. In the 2013 AHS, this is roughly half of the median property tax payments, thus we set ς^I to 0.005.

The fixed refinancing cost ς^r is set to 3,000 in 2013 dollars and is the sum of application, appraisal, inspection, and survey fees, along with attorney review, and title search and insurance costs. Data on the different costs are taken from the Federal Reserve. We use the average of the low and high estimates for these costs.[12]

Labor income

In this section, we outline the central elements of our estimation procedure, and relegate a more detailed description of the data and estimation method to Appendix 1.D. The labor income process is similar to that of Cocco et al. (2005). We estimate a deterministic life-cycle profile of earnings and include the idiosyncratic earnings risk via permanent and transitory shocks. At each age j, household i receives exogenous earnings y_{ij}. For any household, the log earnings before retirement are

$$\log(y_{ij}) = \alpha_i + g(j) + n_{ij} + \nu_{ij} \quad \text{for } j \leq J_{ret}, \tag{1.18}$$

where α_i is a household fixed effect with distribution $N(0, \sigma_\alpha^2)$. The function $g(j)$ represents the hump-shaped life-cycle profile of earnings. The remaining two terms, ν_{ij} and n_{ij}, capture the idiosyncratic earnings risk. The former is an i.i.d. transitory shock with distribution $N(0, \sigma_\nu^2)$. The latter, n_{ij}, allows for households' earnings to permanently deviate

[12] For the estimates of the different costs, see "A consumer's guide to mortgage refinancing", available at https://www.federalreserve.gov/pubs/refinancings/default.htm.

from the deterministic trend, and is assumed to follow a random walk

$$n_{ij} = n_{i,j-1} + \eta_{ij} \quad \text{for } j \leq J_{ret}, \tag{1.19}$$

where η_{ij} is an i.i.d. shock, distributed $N(0, \sigma_\eta^2)$. All shocks are assumed to be uncorrelated with each other. Note that log earnings are represented by the sum of a permanent component, $\log(z_{ij}) = \alpha_i + g(j) + n_{ij}$, and a transitory component ν_{ij}. The permanent earnings state variable in the model is given by z_{ij}.

During retirement there is no earnings risk. Households receive benefits given by

$$\log(y_{ij}) = \min\left(\log(R) + \log(z_{i,J_{ret}}),\ \log(B^{max})\right) \quad \text{for } j \in]J_{ret}, J], \tag{1.20}$$

where R is a common replacement rate for all households and B^{max} is the maximum amount of benefits a household can receive. For simplicity, retirement benefits are a function of permanent earnings in the last period before retirement only.

Equations (1.18) and (1.19) are estimated using PSID data for the survey years 1970 to 1992, following Cocco et al. (2005). The deterministic life-cycle profile $g(j)$ is estimated by regressing log household earnings on dummies for age, marital status, family composition, and education. We control for household fixed effects by running a linear fixed effect regression. A third-order polynomial is fitted to the mean predicted earnings by age.

The variances of the transitory σ_ν^2 and permanent σ_η^2 shocks are estimated in a similar fashion as in Carroll and Samwick (1997). The variance of the fixed effect shock σ_α^2 is identified as the variance of earnings, net of the deterministic trend value in the first period of working life, that is not explained by the estimated variances of the transitory and the permanent shocks. Table 1.2 presents the resulting variances.

The maximum allowable benefit during retirement, B^{max} in equation (1.20), is calculated using data from the Social Security Administration (SSA). The common replacement rate R is set to 50 percent, as in Díaz and Luengo-Prado (2008).

Parameter	Description	Value
σ_α^2	Fixed effect	0.095
σ_η^2	Permanent	0.030
σ_ν^2	Transitory	0.028

Table 1.2: Estimated variances of three-year shocks
Note: During working age, households receive permanent and transitory earnings shocks. In addition, households obtain a fixed effect shock when they enter the economy. During retirement there is no earnings risk. Estimated using PSID data.

1.4.2 Estimated parameters

Table 1.3 shows the structural parameters calibrated by simulated method of moments, along with a comparison between data and model moments. Unless otherwise stated, we use data from the SCF, where we pool the 1989 to 2013 survey years.

Parameter	Description	Value	Target moment	Data	Model
β	Discount factor	0.93	Median LTV	0.35	0.35
δ^r	Depreciation rate, rentals	0.047	Homeownership rate, age < 35	0.44	0.44
$\underline{h}$	Minimum owned house size	137.0	Homeownership rate	0.70	0.70
α	Consumption weight in utility	0.76	Median house value-to-earnings	2.30	2.30
$\bar{q}$	Luxury parameter of bequest	135.6	Net worth p75/p25, age 68-76	5.30	5.61
υ	Utility shifter of bequest	6.5	Median net worth, age 75/50	1.43	1.43
λ	Level parameter, tax system	1.66	Average marginal tax rates	0.13	0.13
SD	Standard deductions	8.02	Itemization rate	0.53	0.53
τ^p	Progressivity parameter	0.14	Distr. of marginal tax rates	See text	

Table 1.3: Estimated parameters
Note: Parameters calibrated by simulated method of moments. The third column shows the resulting parameter values from this estimation procedure. The values are annual when applicable. When simulating the model, we adjust these parameter values to their three-year (one model period) counterparts. The minimum owned house size $\underline{h}$ and the standard deduction SD are in 1000's of 2013 dollars. The fifth column presents the values of data moments that are targeted. The last column shows the model moments that are achieved by using the parameter values in column three.

Although all the parameters are jointly determined in the simulated method of moments, some parameters are especially important for some moments. The discount factor β impacts households' savings and borrowing decisions. Hence, this parameter is used to match the median LTV. The depreciation rate of rental housing δ^r affects how favorable owner-occupied housing is relative to rental housing, which in turn impacts how early in life households become homeowners. Therefore, this parameter is used to target the homeownership rate for those under the age of 35. The minimum owner-occupied house size $\underline{h}$ is calibrated to match the overall homeownership rate. The parameter α determines the weights on consumption and housing services in the utility function. We use this parameter to calibrate the median house value relative to earnings, conditional on owning. The bequest function has two parameters; $\bar{q}$ determines the extent to which bequests are luxury goods, and υ determines the strength of the bequest motive. The former is calibrated to capture the dispersion in net worth among old households, measured as the ratio of net worth in the 75th percentile to the 25th, for ages 68 to 76. The latter is calibrated to fit the difference in net worth between working-age and retired households. As a target, we use the ratio of median net worth for ages 75 and 50. We use the parameter λ, which governs the level of the convex tax and transfer function $T(\tilde{y})$, to target the average marginal tax rate. The target is taken from Harris (2005). We calibrate the standard deduction to match the fraction of the working-age population that itemize tax deductions. Using self-reported rates for working-age households, the itemization rate is 0.53.[13] Our calibrated standard deduction is about $8,000$ in 2013 dollars, which is within the range of standard deductions available to single filers ($6,100) and married households filing jointly ($12,200) in 2013.

The parameter determining the progressivity of the federal labor income tax τ^p, is set to match the distribution of households exposed to the different statutory marginal tax rates. We minimize the sum of the absolute values of the differences between the shares of households exposed to the statutory tax brackets in data compared to in the model. For this estimation procedure, we allocate households to their nearest tax bracket in the model, and we use data on shares from the Congressional Budget Office in 2005 (Harris, 2005). The statutory tax brackets we use are consistent with the tax code from 2003 to 2012 (The Tax Foundation, 2013). The resulting progressivity parameter value is 0.14, which is close to that in Heathcote et al. (2017). Figure 1.1 displays the fractions of the working-age population exposed to the different statutory marginal tax rates in the data (Harris, 2005) versus in the model.

[13] In this case, we do not include households aged 23-25 when we compute the model moment. These ages correspond to the first model period, where households by construction cannot deduct property taxes or mortgage interest payments. Hence, the itemization rate is artificially low in the model for this age group.

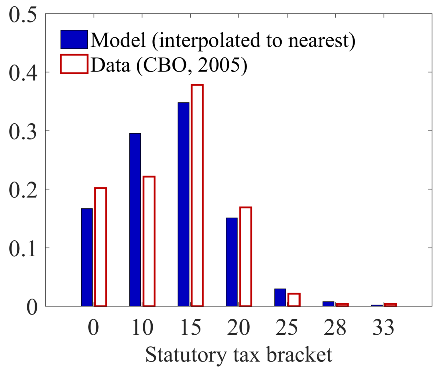

Figure 1.1: Fractions of taxpayers facing different marginal tax rates
Note: The model refers to the results from the initial steady state with MID. For comparison purposes, we interpolate households' marginal tax rates to the nearest tax brackets, as the labor income tax schedule is continuous in the model. The data is from Harris (2005).

1.4.3 Model fit

As is evident in Table 1.3, the calibration enables the model to successfully hit the target moments. However, the reliability of our results does not only depend on how well the model performs with respect to aggregate measures. It also depends on the distributions and life-cycle profiles of relevant variables.

The life-cycle profiles of homeownership, LTV, and mortgage-to-earnings are key indicators of the heterogeneity in exposure to the mortgage interest deductibility. Comparisons to SCF are displayed in Figure 1.2. The model performs well with respect to these variables, both in terms of magnitudes and life-cycle patterns, although there are some discrepancies. The model also produces a decent fit of the median house-to-earnings, which is a measure of exposure to price changes in the housing market. The jump in the median house-to-earnings at age 65 in the model is a result of households retiring with certainty at that age.

Data on U.S. tax returns and the SCF show that the fraction of households that itemize deductions is increasing in earnings and that there is a strong skewness in MID claims.[14] In the 2013 tax filings, only about

[14]The tax return data is publicly available at the IRS webpage. We use data from "SOI tax stats - individual statistical tables by size of adjusted gross income", tables 1.4 and 2.1.

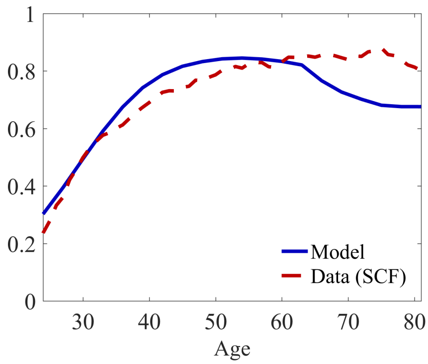

(a) Homeownership rate

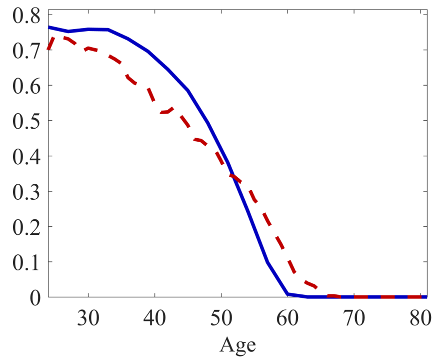

(b) Median LTV

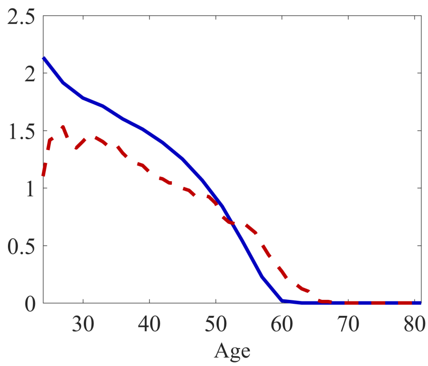

(c) Median mortgage-to-earnings

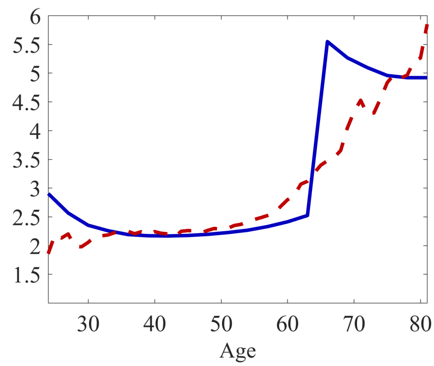

(d) Median house-to-earnings

Figure 1.2: Comparison of model versus data: non-targeted profiles
Note: The model refers to the results from the initial steady state with MID. The data is taken from Survey of Consumer Finances (SCF), survey years 1989-2013.

four percent of those earning less than $15,000 (24 percent of all returns) itemized deductions, and they merely claimed two percent of all mortgage interest deductions. This stands in sharp contrast to comparable numbers for those earning more than $100,000 (top 15 percent). They claimed 55 percent of the total mortgage interest deductions, and more than 82 percent used itemized deductions. A similar skewness is apparent in the SCF, although somewhat less pronounced. As seen in Figure 1.3a and Figure 1.3b, our model is able to replicate these important patterns: high earners itemize the most and claim a disproportionately large share of the mortgage interest deductions.

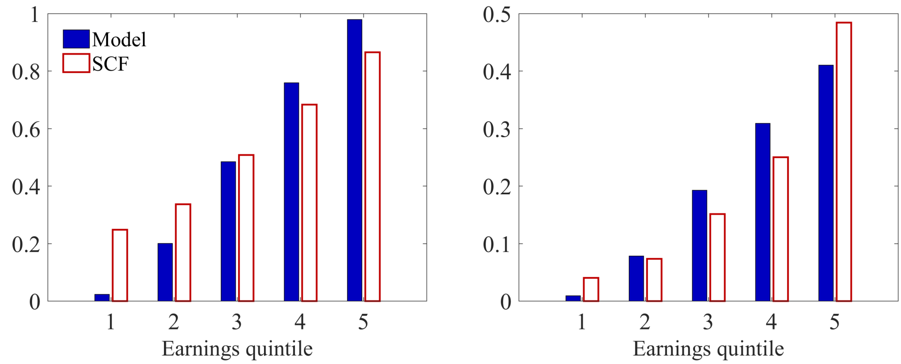

(a) Fractions that itemize

(b) Fractions of mortgage interest deductions

Figure 1.3: Itemizers and MID claims in the initial steady state, across earnings quintiles

Note: Working-age households only. The data is taken from the SCF, survey years 1995-2013 (the data on itemization is missing in the 1989 and 1992 waves). Mortgage interest deductions are computed from reported mortgages and interest rates for those who itemize.

1.5 Results

1.5.1 What are the long-run effects of removing the MID?

What would the level of house prices in the U.S. be if households were not able to deduct mortgage interest payments? Does the MID promote homeownership? What fraction of American households would prefer to be born into a world without the MID, and how much would they gain or lose?

These questions regarding the long-run implications of removing the MID are all addressed in this section. Although the focus of this paper is on the transitional dynamics of repealing the MID, the answers to these questions are also relevant for our purpose. Indeed, it is difficult to motivate a study of the short-run dynamics if the long-run welfare effects are negative. Moreover, the key mechanisms in the long run are also at work in a transition.

In order to study the long-run effects of removing the MID, we compare the initial steady state with MID to a new steady state in which the possibility to deduct mortgage interest payments is repealed. Specifically, we study the effects of changing the deductibility parameter τ^m from the initial value of one to zero, while imposing tax neutrality and accounting for changes in bequests. The labor income tax level parameter λ is adjusted so that the net tax revenue for the government is unchanged between the steady states. We alter the bequest parameter γ to distribute any changes in bequests.

Prices and aggregates

Table 1.4 presents a comparison of the two steady states for a number of key variables. Overall, the new steady state without MID is characterized by lower house and rental prices, higher homeownership, reduced indebtedness, lower taxes, and more bequests. The price decrease is driven by a downward shift in the demand for housing among homeowners who often itemize. These households experience an increase in the user cost of owning, as discussed in Section 1.3. If the house price is held constant, households in this group would wait longer until they buy their first house, and buy smaller houses. When the house price is allowed to decline, households who often itemize do no longer postpone their house purchases, but they still demand smaller houses. Overall, in the new steady-state equilibrium, the homeownership is virtually unchanged for this group of households, whereas they demand smaller houses.

For those who seldom itemize, the lower house price has a positive effect on homeownership. Some households who would never own a house in the initial steady state are homeowners in the new steady state. Indeed, the fraction of households that own a house at some point in life increases by about one percentage point (see *fraction ever-owner* in Table 1.4). Moreover, those who own a house but seldom itemize in the initial steady state choose to buy their first house earlier in the new steady state. Overall, the homeownership rate increases by approximately one percentage point to around 71 percent. This result confirms the findings and the underlying mechanism in Sommer and Sullivan (2018). They document that removing the MID is associated with an increase in the homeownership rate due to the fall in the house price.

In Table 1.4, we see that the mean mortgage level decreases significantly. This is primarily driven by households that often itemize. The fall in the mortgage level is not only caused by the higher cost of mortgage financing, but also by the change in house sizes and the fall in the house price. Since it is the itemizing households that demand smaller houses and are directly affected by the MID, they are also those that decrease their mortgage levels the most.

	MID	No MID	Relative difference (%)
House price	1	0.965	-3.47
Rental price	0.238	0.234	-1.66
Homeownership rate	0.70	0.71	1.88
Fraction ever-owner	0.88	0.89	1.59
Mean owned house size	215	211	-2.15
Mean LTV	0.36	0.31	-12.09
Mean mortgage	74	60	-19.29
Mean bond holdings	20.6	21	1.81
Mean marginal tax rate	0.150	0.146	-2.59
Mean bequest collected	152	158	3.57

Table 1.4: Long-run effects on prices and aggregates of removing the MID
Note: The first column shows prices and aggregate measures in the initial steady state with MID, whereas the second column shows the corresponding values in the steady state without MID. The rental price corresponds to a three-year (one model period) cost of renting. "Fraction ever-owner" is the fraction of households that own a house at some point during their life. The mean house size, LTV, and the mortgage level are conditional on owning. The mean owned house size, mortgage, bond holdings, and bequest collected are in 1000's of 2013 dollars. The mean marginal tax rate is gross of deductions.

Why are U.S. households better off without the MID in the long run?

We use the ex-post consumption equivalent variation (CEV) as our welfare measure. This is defined as the per-period percentage change in realized consumption that is required in the steady state with MID to make a household indifferent to an economy without MID. Formally, let $\tilde{V}$ be the discounted welfare and $(\tilde{c}_{i,j}, \tilde{s}_{i,j}, \tilde{q}'_{i,j})$ be the realized consumption, housing services, and net worth in the steady state without MID,

$$\tilde{V} \equiv \sum_{j=1}^{J} \left(\beta^{j-1} \prod_{k=1}^{j-1} \phi_k \right) \left[U_j(\tilde{c}_{i,j}, \tilde{s}_{i,j}) + \beta(1-\phi_j) U^B(\tilde{q}'_{i,j}) \right].$$

Then, for each household we solve for Δ that makes the discounted welfare under the two tax regimes equal

$$\sum_{j=1}^{J}\left(\beta^{j-1}\prod_{k=1}^{j-1}\phi_k\right)\left[U_j\left((1+\Delta)c_{i,j},s_{i,j}\right)+\beta(1-\phi_j)U^B(q'_{i,j})\right]=\tilde{V},$$

where $(c_{i,j}, s_{i,j}, q'_{i,j})$ are the realizations of each variable in the steady state with MID. If we set Δ to zero, the left-hand side of this equation is simply the discounted welfare in the initial steady state. In the remainder of the paper we will refer to Δ as CEV. We are also interested in what fraction of households that benefit from a removal, which we define as the share of households with a CEV greater than or equal to zero.

An overwhelmingly large fraction, 88 percent of households, are better off being born into the steady state without MID, see the last column in Table 1.5. On average, the welfare gain of being born into the steady state without MID is equivalent to that of increasing consumption by 0.91 percent in all periods in the initial steady state.

The direct effect of no longer having the mortgage subsidy is negative as seen in the first column of Table 1.5. Yet a large fraction of households experience a small or no loss. Even with MID in place, many households do not itemize deductions. In addition, as seen in Figure 1.3b, the amounts of mortgage interest deductions are highly skewed. Households with higher earnings claim a disproportionately large share of the total mortgage interest deductions. Most itemizing households deduct relatively modest amounts of mortgage interest payments.

There are three equilibrium effects that improve households' welfare: the lower house and rental prices, the lower labor income taxes, and the increased bequests. The lower house price in the steady state without MID makes both rental and owner-occupied housing more affordable, which increases welfare. Importantly, the lower house price reduces the equity requirement and makes the PTI requirement less stringent. This enables more households to become homeowners and allows some households to purchase a house earlier. In the second column in Table 1.5, we see that the price adjustment is sufficient to create significant positive welfare effects, and 74 percent would prefer to live in a world without MID. The lower labor income tax and the increased bequests in the new steady state further increase the welfare for all households. Households at the

Mean CEV (%)	-0.54	0.32	0.78	0.91
Fraction in favor	0.15	0.74	0.85	0.88
Rental and house prices adjust	-	✓	✓	✓
Tax neutrality	-	-	✓	✓
Bequests adjust	-	-	-	✓

Table 1.5: Long-run welfare effects of removing the MID, for newborns
Note: Mean CEV (%) refers to the average consumption equivalent variation in percent, for newborns. For example, the average welfare effect of removing the MID when house prices, taxes, and bequests adjust is equivalent to a 0.91 percent increase in consumption in all periods, in the initial steady state. The fraction in favor is the fraction of newborns with a CEV greater than or equal to zero.

top of the earnings distribution constitute the only group for which the direct negative effect of removing the MID outweighs the benefits of lower equilibrium prices and taxes and higher bequests.

1.5.2 What are the effects of an immediate removal of the MID?

Our results in the previous section suggest that U.S. households would be considerably better off in a world in which they cannot deduct mortgage interest payments. However, the long-run analysis does not touch upon another important question: is a repeal of the MID also beneficial for current households? To shed some light on this question, we need to consider the short-run effects. In this section, we present the results of an immediate removal of the MID, i.e., $\tau_t^m = 0$ for all $t \geq 1$.

Who are the winners and losers from a reform?

In order to evaluate the welfare effects of the immediate removal, we consider the lifetime gains and losses incurred by households alive when the policy is implemented. These welfare effects can differ markedly from the long-run analysis, as households have made long-term decisions based on the presumption that they can deduct mortgage interest payments. As is shown in the analysis below, the welfare effects therefore tend to be significantly lower and much more dispersed.

Similar to the steady-state analysis, there are four main factors influencing how a household is affected by the removal policy. First, households that itemize deductions and have a mortgage are directly negatively affected by a reduction of the MID. Second, the sudden fall in the house price, as seen in Figure 1.4, reduces the home equity for existing homeowners, while renters benefit from less stringent constraints in the housing market and lower rental prices. Third, all households are positively affected by an instant fall in the labor income tax level since the government no longer subsidizes mortgage financing. Finally, the unexpected fall in the house price lowers the values of bequests, which affects all households negatively. A detailed overview of the transitional dynamics is presented in Section 1.5.3, where we compare the immediate policy to alternative removal policies.

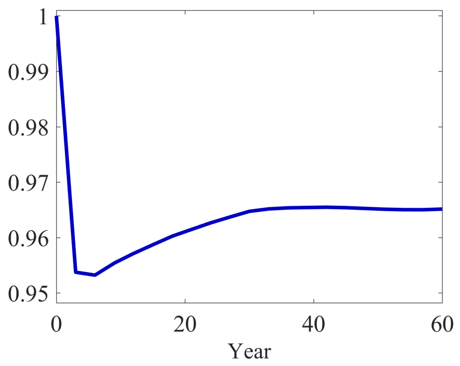

Figure 1.4: House price dynamics from an immediate removal of the MID

On average, the immediate removal policy results in significant welfare losses for current U.S. households. The average welfare effect of an immediate removal is equivalent to a 0.4 percent decrease in consumption in all remaining periods in the initial steady state and only 39 percent would gain from such a reform. This stands in sharp contrast to the long-run analysis, where 88 percent would benefit from a world without MID.

Furthermore, we could not find a one-time cash transfer scheme that, in combination with the removal, would lead to a Pareto improvement. Taxing all winners and compensating all losers such that every household is indifferent between a reform and no reform would produce a transfer-scheme deficit of 1.2 percent of average cash-on-hand.

The aggregate results mask important heterogeneous welfare effects. Figure 1.5 displays the distribution of welfare changes in the first period of the transition. Based on this distribution, we allocate households into four groups as indicated by the vertical lines in the figure. The first group contains the households who experience the largest welfare losses in the transition. The second group contains the households with less extreme, but still sizable negative CEVs. The third group is made up by a mass of households that have CEVs around zero. The households in the right bell of the distribution are allocated to the fourth group. Table 1.6 presents key characteristics for the different groups.

The bimodal shape of the CEV distribution stems from differences in welfare effects between homeowners and renters. The mass around the right-hand peak consists of renters, while the mass around the left-hand peak, groups one to three, consists of homeowners. Renters are not directly affected by the removal of the MID, but benefit from the lower rental price, relaxed LTV and PTI constraints in the housing market, and lower taxes.

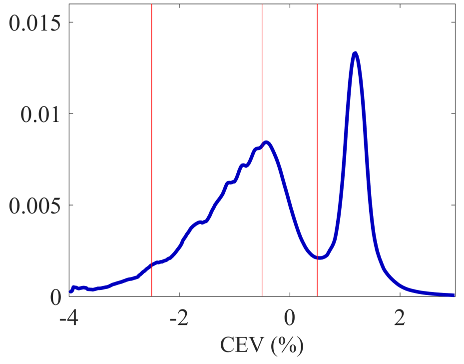

Figure 1.5: The distribution of welfare effects under the immediate removal policy

Note: CEV (%) refers to the ex-post consumption equivalent variation in percent, for all households that are alive in the first period of the transition. The vertical lines allocate households into different groups based on their welfare effects. See Table 1.6 for key characteristics of these groups.

Group:	1	2	3	4
CEV range:	< -2.5	$[-2.5, -0.5[$	$[-0.5, 0.5[$	≥ 0.5
Working age:				
Mean CEV	-3.09	-1.18	-0.13	1.23
Homeownership rate	1	0.98	0.87	0.03
Itemization rate	0.99	0.84	0.56	0.01
Age	38	45	47	37
Earnings	133	106	85	44
House size	310	231	180	165
Mortgage	219	119	67	71
LTV	0.71	0.54	0.41	0.42
Retirement age:				
Mean CEV	-4.57	-1.39	-0.19	1.19
Homeownership rate	1	1	1	0.03
Age	79	70	68	74
Earnings	42	40	28	15
House size	235	218	161	146
LTV	0.02	0.03	0.05	0.07

Table 1.6: Characteristics of winners and losers in the short run
Note: Groups 1 to 4 correspond to the four groups indicated by the vertical lines in Figure 1.5. Thus, the welfare effects are those experienced under the immediate removal policy. Other measures correspond to mean values of households in the event that the MID is not repealed. House size, mortgage, and LTV are conditional on owning a house. Earnings, house size, and mortgage are in 1000's of 2013 dollars.

Homeowners realize several negative effects in the short run, but the extent to which they are affected varies with the household characteristics. By comparing the three groups of homeowners in Table 1.6, we see that the CEV is decreasing in mortgages, permanent earnings, and the itemization rate for working-age households. Homeowners with larger mortgages and higher earnings benefit more from itemizing deductions. Consequently, they are relatively worse off when they can no longer deduct mortgage interest payments, as represented by the long, thick tail of negative values in Figure 1.5. Table 1.6 also shows that households with lower CEVs on average own larger houses, which primarily reflects that these households are high earners. In addition, younger homeowners tend to be worse off. This mainly follows from younger homeowners having higher LTVs. For retired households, the very old with large houses are the biggest losers. These households rely heavily on housing equity as they have low earnings relative to wealth, and thus suffer significantly from the house price fall.

The transition also entails positive effects for homeowners, although these are generally outweighed by the negative effects. All homeowners benefit from the lower labor income taxes when the MID is removed, as well as the decreased property tax payments following the fall in the house price.

The results in Table 1.6 also help explain why a one-time cash transfer between households cannot achieve a Pareto improvement. Not only is a majority of households against the reform, but those who are hurt by the removal tend to have higher life-time earnings. This negative correlation between income and welfare implies that relatively large transfers in absolute terms are required to compensate the losers.

Why do we find negative welfare effects?

Other papers that have studied the short-run welfare effects of removing the MID find that a majority of households would benefit from an immediate policy; see Floetotto et al. (2016) and Sommer and Sullivan (2018). Our model differs from those in earlier papers along a variety of dimensions. Although our model does not nest theirs, three major modeling features account for most of the differences in the welfare effects relative to Sommer and Sullivan (2018), which is arguably the paper closest to ours in terms of modeling. These features include having a refinancing cost of mortgages, an explicit modeling of households in retirement, and accounting for changes in bequests caused by a lower house price in the transition.

A refinancing cost of mortgages makes housing equity less liquid and it is more difficult for homeowners to cushion negative shocks. The refinancing cost makes it costly to increase a mortgage, and it is only worthwhile to increase a mortgage by a relatively large amount. In the transition, the house price decline limits the amount by which households can increase their mortgage, through the LTV constraint. Thus, households who receive negative shocks during the transition are less inclined to use housing equity to smooth their consumption, resulting in lower welfare. Furthermore, refinancing costs lead to a larger house price decline early in the transition, through a similar mechanism. In an economy where mortgage refinancing is costly, households are more inclined to get access to their housing equity by selling their home, which has a negative impact on the aggregate price level.

As we explicitly model the full life-cycle of households, we are able to study the welfare effects of retirees. We find that homeowners in retirement are relatively worse off when the MID is removed. Older households hold more housing wealth, and their housing wealth relative to earnings is substantially higher as depicted in Figure 1.2d. Furthermore, older households have fewer periods left to smooth the negative wealth shock resulting from the house price drop. Finally, because retirees have a higher probability of dying, they care more about the bequests they leave behind. Thus, for many retirees, the unexpected fall in net worth in the transition lowers their welfare.

In our analysis, households are also negatively affected by a reduction in the values of bequests received. This is crucial at the beginning of the transition when the house price fall sharply reduces wealth. When households receive less bequests, there is a reduction in welfare. In contrast, Sommer and Sullivan (2018) use a standard assumption that any accidental bequests are fully taxed and that the government spends this revenue on activities that do not affect any agents in the economy.

In a calibration where we remove the cost of refinancing, consider the welfare effects of the working-age population only, and do not distribute changes in bequests, we find a positive average welfare effect that is more in line with previous studies. In the first column of Table 1.7, we can see that under these assumptions, the average CEV is 0.03 percent and a majority (52 percent) of households are in favor of an immediate removal of the MID. Moreover, the results in this table suggest that all three model features are central for understanding why our welfare results are lower than those in Sommer and Sullivan (2018).

1.5.3 Do households prefer more gradual removal policies?

In the previous section, we show that an immediate removal of the MID results in considerable negative welfare effects on average. The negative effects are primarily driven by homeowners who can no longer deduct mortgage interest payments and who suffer from losses in their housing equity. Given the large positive long-run welfare effects of an MID removal, an investigation of alternative removal policies that could potentially improve the welfare effects for homeowners is warranted.

To enable homeowners to adjust their asset allocations before the MID is repealed, we consider two policies in which the MID is removed

Mean CEV (%)	0.03	-0.14	-0.22	-0.40
Fraction in favor	0.52	0.46	0.41	0.39
Include welfare retirees	-	✓	✓	✓
Bequests adjust	-	-	✓	✓
Refinancing cost	-	-	-	✓

Table 1.7: Model features that can explain our negative welfare result
Note: Welfare results of an immediate removal policy. The first column shows the results from a model specification where we do not: i) include the welfare effects of retirees; ii) adjust bequests received by households; and iii) include refinancing costs of mortgages. The last column shows our benchmark result. The fraction in favor is the fraction of households with a CEV greater than or equal to zero. For a description of CEV (%) see *Note* below Figure 1.5.

less rapidly.[15] Under a *gradual* policy, households can deduct mortgage interest payments for another 15 years (5 model periods), but the deductible share is reduced stepwise over that period such that $\{\tau_t^m\}_{t=1}^{t=\infty} = \{1, 0.8, 0.6, 0.4, 0.2, 0, 0, ...\}$. We also study an *announcement* policy in which households are informed that all interest payments can be deducted for another 15 years before the MID is removed permanently, i.e., $\{\tau_t^m\}_{t=1}^{t=\infty} = \{1, 1, 1, 1, 1, 0, 0, ...\}$. These policies together with the immediate reform are depicted in Figure 1.6a.

How do short-run dynamics depend on the timing of policies?

Figure 1.6 shows the short-run dynamics for the house price, the rental price, the average marginal labor income tax rate before deductions, and the bequest rate, for all three policies. The house price falls most rapidly under the immediate policy. The price fall under a given removal policy is mainly driven by changes in the housing demand of young itemizing households. As seen in the second row of Figure 1.2, young households have high LTVs and mortgage-to-earnings when they enter the housing market. As these households also tend to itemize deductions, they respond strongly to changes in the deductibility rate. Under the gradual and announcement policies, the response in housing demand is smaller due to the extended possibilities to deduct mortgage interest payments.

[15] For an analysis of a grandfather policy, see Appendix 1.E.

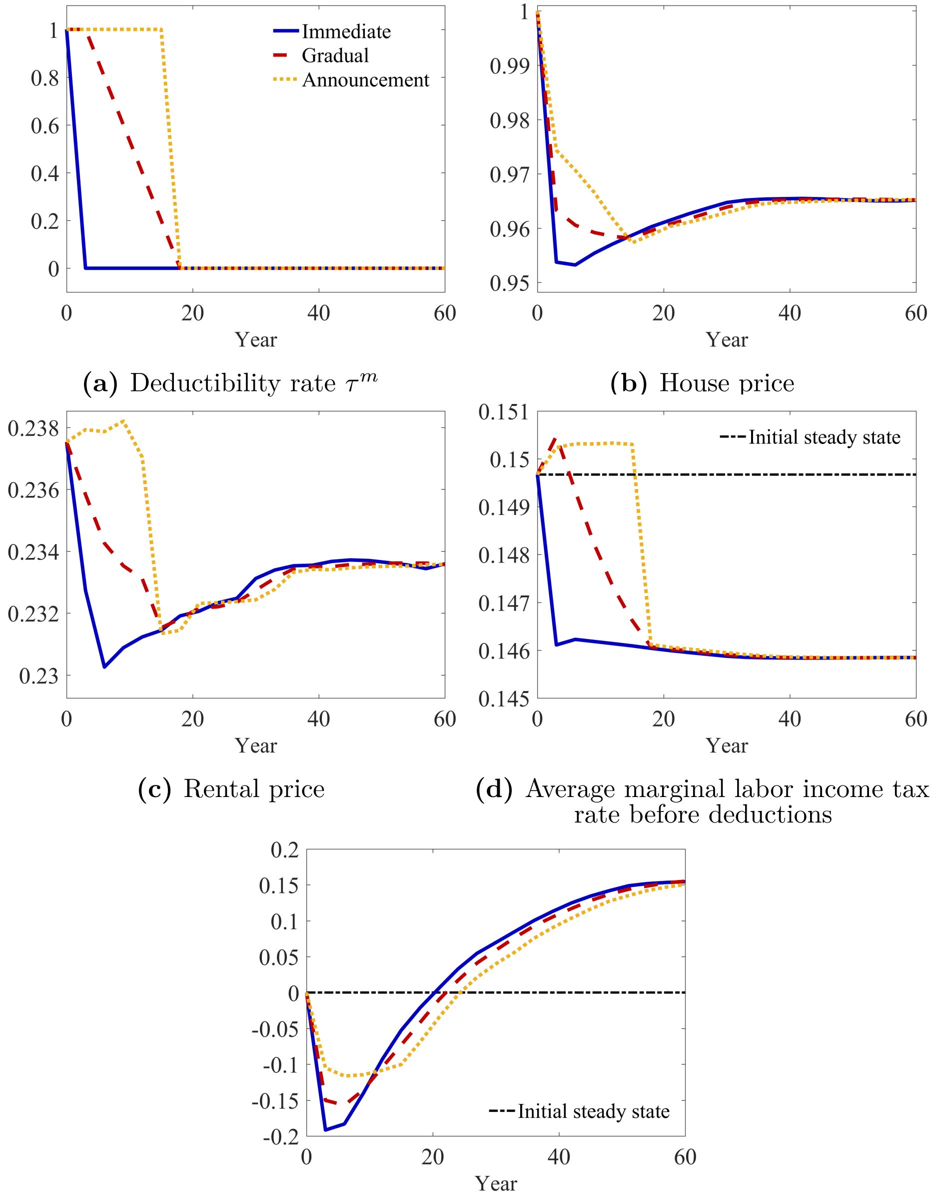

(a) Deductibility rate τ^m

(b) House price

(c) Rental price

(d) Average marginal labor income tax rate before deductions

(e) Bequest rate γ (%)

Figure 1.6: Short-run dynamics from removing the MID, across policies

Note: Panel (a) shows how the deductibility rate is decreased under the three policy reforms. All policies are implemented unexpectedly and households have perfect foresight of the transition paths of prices, taxes, and bequests. Panels (b)-(e) show how the house price, the rental price, the average marginal tax rate before deductions, and the bequest rate behave in the short run, in response to the paths of the deductibility rate. The rental price corresponds to a three-year (one model period) cost of renting.

Although the instantaneous drop in the house price is the largest under the immediate policy, more than fifty percent of the total price fall occurs in the first transition period for the gradual and announcement policies. The higher present value of the future user cost of owning instantly results in a lower demand for owned housing, under all policies. The demand effect is reinforced by the transaction costs associated with buying and selling a house. The transaction costs restrain households from frequently re-optimizing their house size, which makes a house purchase a long-term investment.

The short-run dynamics of the rental price is fully explained by the path of house prices, as shown in equation 1.9. Under the immediate policy, the rental price closely follows the house price, whereas the rental price remains elevated for some periods under the more gradual policies.

The differences in the labor income tax levels between policies are driven by the paths of the deductibility rate and the house price. A lower mortgage deductibility rate decreases the government's tax expenditures, and allows the government to reduce the labor income tax level. On the other hand, a fall in the house price decreases the property tax payments, which worsens the government's budget. Under the immediate policy, the labor income tax level can be reduced at once. Under the gradual and announcement policies, the labor income tax rates initially increase, as the revenue from property taxes falls and the government still spends large amounts on interest deductions.

The initial drop in the bequest rate is driven by the unexpected fall in the house price, which decreases the value of collected bequests. As the fall in the house price is the largest under the immediate policy, so is the negative effect on bequests. Along the transition, the bequest rate increases as households' asset holdings become increasingly similar to those in the new steady state, where the average net worth is higher.

How do welfare effects depend on the timing of policies?

Although the less abrupt policies give households more time to adjust their allocations, we find that the immediate policy is actually preferred to those policies. As seen in Table 1.8, the immediate policy has both the highest mean CEV and is the policy with the highest share of households experiencing welfare improvements. Thus, we find that none of the policies are able to achieve a majority support or positive welfare effects on average. Even in an analysis where we consider the discounted welfare of all households that enter the economy along the transition, the average welfare effect remains negative for all policies.[16] Moreover, we cannot find a one-time cash transfer scheme that results in a Pareto improvement under any of the policies considered.

	Immediate	Gradual	Announcement
Mean CEV (%)	-0.40	-0.52	-0.52
Fraction in favor	0.39	0.35	0.27

Table 1.8: Welfare effects of households alive in the first period of the transition
Note: The fraction in favor is the fraction of households with a CEV greater than or equal to zero. For a description of CEV (%) see *Note* below Figure 1.5.

There are substantial differences in the CEV distributions across policies, as seen in Figure 1.7. Naturally, the direct effect of removing the MID is negative under all policies. The average welfare loss from this channel is dampened under the gradual and announcement reforms, which reduces the left-hand tail of the CEV distribution.

The slower fall in rental prices and house prices under the gradual and announcement policies affects both renters and homeowners. Renters prefer the immediate policy, since they benefit from a faster decline in prices. As a result, the right-hand peak of the distribution in Figure 1.7 is shifted to the left under the gradual and announcement policies. For homeowners, the accelerated fall in the house price under the immediate policy reduces the housing equity more rapidly, and the losses distributed from the rental companies are higher. The overall effect of changes in rental prices and house prices is a decrease in average welfare. Quantitatively, this negative effect is similar in magnitude under all policies.[17]

The fall in house prices also leads to a reduction in bequests during the first periods of the transition and has a negative impact on all households. This negative effect is somewhat less pronounced under the more gradual policies when the house price fall is smaller.

[16] Specifically, the mean discounted CEV (%) would be -0.08, -0.14, and -0.16 under the immediate, gradual, and announcement policy, respectively. We discount the welfare of newborns by β^{t-1}, noting that $t = 1$ is the first period of the transition.

[17] For a detailed account of the welfare effects under different equilibrium assumptions, see Figure 1.10 in Appendix 1.F.

A lower labor income tax level benefits all households and shifts the whole CEV distribution to the right. Households benefit the most from labor income tax changes under the immediate policy, which has the lowest tax rate in the first five periods of the transition. The short-term differences in tax rates between policies have important implications for welfare and constitute a key reason why the immediate policy achieves the smallest welfare loss.

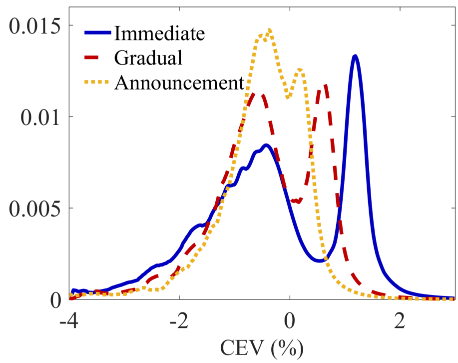

Figure 1.7: Distributions of short-run welfare effects, across policies
Note: Distributions of welfare effects of the three policies, for households alive in the first period of the transition. For a description of CEV (%) see *Note* below Figure 1.5.

1.5.4 An MID removal after the Tax Cuts and Jobs Act

At the end of 2017, the Tax Cuts and Jobs Act (TCJA) was enacted (see, e.g., Gale et al. (2019) for a summary). In this section, we take a closer look at the welfare effects of an MID removal after incorporating some of the main changes of the tax reform. Specifically, we focus on two changes to the tax system: the near doubling of the standard deduction and the new cap on deductions for state and local income taxes and property taxes. These changes are likely to be particularly important for an MID removal. They reduce the fraction of households that choose to itemize deductions and thus the number of households that benefit from the MID. There are other features of the fiscal reform that we have not incorporated in the model because we believe that they are unlikely to have large effects on our results.[18]

We operationalize the TCJA by increasing the baseline standard deduction by a factor of 1.9 and by setting the maximum deduction for the sum of state and local income taxes and property taxes to 10,000 in 2018 dollars.[19] For simplicity, we assume that the new legislation is permanent, although these individual tax code provisions are all scheduled to expire at the end of 2025. Note that we do not require the TCJA to be tax neutral, i.e., the labor income tax level is not changed. However, we do adjust the bequest parameter γ, taking into account that the bequests left behind may change. We proceed by repeating the policy experiments in the previous section, but take as a starting point the steady state with taxes set according to the TCJA.

Table 1.9 summarizes the results of the short-run policy experiments, whereas the long-run results are provided in Appendix 1.G. For all removal policies, a majority of households are against a removal and the average CEV is negative. Quantitatively, the average welfare effects are less negative compared to our benchmark results, as the direct effect of removing the MID is reduced under the new tax code. Under the TCJA tax code, only households with considerable mortgages find it worthwhile to itemize tax deductions, resulting in an itemization rate of just 9 percent. Since

[18]There are primarily three parts of the tax reform that are related to our modeling framework that we have chosen to not incorporate. First, under the TCJA it is no longer possible to deduct interest payments for home equity lines of credit. We have no explicit role for home equity lines of credit in the model and only 5 percent of total mortgages are home equity loans in the SCF 2013 wave. Second, the cap on total mortgage interest payments that can be deducted was reduced from 1M to 750k. In our model, this change affects very few households, especially since the new cap on property tax deductions reduces the house sizes of high-income households. Finally, the TCJA reduced the tax rates and altered the thresholds for most federal income tax brackets. In the model, we calibrate the two parameters of our labor income tax function to match the average marginal tax rate in data, and the distribution of households exposed to the different statutory marginal tax rates. We do not have data for this after the new tax rates and thresholds were implemented, and it is therefore not obvious how the changes should be translated into changes of the parameters. However, with lower marginal tax rates for high-income households, the benefits of the MID are likely further reduced with the new tax schedule. As a result, the negative effects of a removal may be smaller.

[19]Under prior law, the 2018 standard deduction would have been 6,500 dollars for single filers, 13,000 dollars for joint filers, and 9,550 dollars for head of household. Under the TCJA, the standard deduction is 12,000 dollars for single filers, 24,000 dollars for joint filers, and 18,000 dollars for head of household; see Gale et al. (2019).

removing the MID affects fewer households directly, the removal also has a more muted effect on taxes and prices. For example, the house price fall is only about half as large as under the baseline calibration. As a result, the welfare losses for homeowners are smaller, but so are the welfare gains for renters.

	Immediate	Gradual	Announcement
Mean CEV (%)	-0.28	-0.30	-0.26
Fraction in favor	0.39	0.38	0.35

Table 1.9: Short-run welfare effects: Tax Cuts and Jobs Act
Note: The fraction in favor is the fraction of households with a CEV greater than or equal to zero. For a description of CEV (%) see *Note* below Figure 1.5.

1.6 Concluding remarks

A growing academic literature consistently shows that, in the long run, most American households would be better off without the MID. Much less is known about how a repeal of the MID would affect current households and, in particular, how these effects depend on the design of the removal policy. In this paper, we attempt to fill this gap by taking into account transitional dynamics and studying the welfare effects of several MID removal policies.

Our results show i) that the welfare effects of an unexpected and immediate removal policy are negative on average and less than forty percent of households benefit from the reform, and ii) that more gradual policies do not improve these outcomes. The results materialize despite our finding that 88 percent of households would prefer to be born into a world without the MID. We argue that the inclusion of mortgage-refinancing costs, which reduce the liquidity of housing wealth, and an explicit modeling of retirees, are the main reasons why we find considerably lower welfare effects as compared to the existing literature. In our analysis, we find that both aggregate and distributional welfare measures depend significantly on how the MID is removed and that households differ in their preferred policy design. More gradual policies, which give households more time to prepare for an MID removal, are successful in mitigating the losses for those who suffer the most under an immediate policy. However, a majority of households actually prefer an immediate removal with large and instantaneous equilibrium effects of lower prices and taxes.

Our analysis highlights the importance of including realistic life-cycle dynamics and key frictions to understand the welfare effects of tax policies in the housing market. To further increase our comprehension of how government policies affect households differentially, this class of heterogeneous agent models provide a promising platform. There are a number of extensions that are worthwhile considering in future work on housing tax reforms, in particular when studying a removal of the MID. First, potential demand effects on output from, e.g., lower house prices could be explored. To the extent that such changes in output can have important feedback effects into house prices, these effects are omitted from our analysis. Second, it would be interesting to explore whether a Pareto improvement can be achieved by combining the removal with more elaborate transfer schemes. In this paper, we do not find a one-time transfer scheme between winners and losers of the current generation that would make everyone better off. However, since future generations benefit from the removal, it might be possible to obtain a Pareto improvement by allowing the government to take up debt and redistribute gains from coming generations. Last, expanding the analysis by allowing house prices to be non-linear in house size may have implications for homeownership and welfare. Our analysis shows that a removal of the MID reduces the demand for larger houses, whereas more households buy smaller homes. Although we find these considerations interesting, we leave them as suggested avenues for future research.

References

Bell, F. C. and M. L. Miller (2005, August). Life tables for the United States Social Security Area 1900-2100. Actuarial study No. 120, Social Security Administration.

Boar, C., D. Gorea, and V. Midrigan (2020). Liquidity constraints in the U.S. housing market. *Working Paper*.

Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters 91*(3), 312–320.

Carroll, C. D. and A. A. Samwick (1997). The nature of precautionary wealth. *Journal of Monetary Economics 40*(1), 41–71.

Chambers, M., C. Garriga, and D. E. Schlagenhauf (2009). Housing policy and the progressivity of income taxation. *Journal of Monetary Economics 56*(8), 1116–1134.

Cho, S.-W. and J. Francis (2011). Tax treatment of owner occupied housing and wealth inequality. *Journal of Macroeconomics 33*, 42–60.

Cocco, J. F., F. J. Gomes, and P. J. Maenhout (2005). Consumption and portfolio choice over the life cycle. *Review of Financial Studies 18*(2), 491–533.

Conesa, J. C. and D. Krueger (1999). Social security reform with heterogeneous agents. *Review of Economic Dynamics 2*(4), 757–795.

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

Desmond, M. (2017). How homeownership became the engine of American inequality. *The New York Times Magazine*. May 9th, 2017.

Díaz, A. and M. J. Luengo-Prado (2008). On the user cost and homeownership. *Review of Economic Dynamics 11*(3), 584–613.

Druedahl, J. and T. H. Jørgensen (2017). A general endogenous grid method for multi-dimensional models with non-convexities and constraints. *Journal of Economic Dynamics and Control 74*, 87–107.

Floetotto, M., M. Kirker, and J. Stroebel (2016). Government intervention in the housing market: who wins, who loses? *Journal of Monetary Economics 80*, 106–123.

Gale, W. G., H. Gelfond, A. Krupkin, M. J. Mazur, and E. J. Toder (2019). Effects of the Tax Cuts and Jobs Act: a preliminary analysis. *National Tax Journal 71*(4), 589–612.

Gervais, M. (2002). Housing taxation and capital accumulation. *Journal of Monetary Economics 49*(7), 1461–1489.

Greenwald, D. (2018). The mortgage credit channel of macroeconomic transmission. *Working Paper*.

Gruber, J. W. and R. F. Martin (2003). Precautionary savings and the wealth distribution with illiquid durables. *International Finance Discussion Papers 773*.

Guvenen, F. (2009). An empirical investigation of labor income processes. *Review of Economic Dynamics 12*(1), 58–79.

Harding, J. P., S. S. Rosenthal, and C. Sirmans (2007). Depreciation of housing capital, maintenance, and house price inflation: estimates from a repeat sales model. *Journal of Urban Economics 61*(2), 193–217.

Harris, E. (2005, November). Effective marginal tax rates on labor income. A CBO paper, Congressional Budget Office.

Heathcote, J., K. Storesletten, and G. L. Violante (2017). Optimal tax progressivity: an analytical framework. *Quarterly Journal of Economics 132*(4), 1693–1754.

Henderson, J. V. and Y. M. Ioannides (1983). A model of housing tenure choice. *American Economic Review 73*(1), 98–113.

JCT (2017, January). Estimates of federal tax expenditures for fiscal years 2016-2020. Joint committee print, Joint Committee on Taxation.

Kaplan, G., K. Mitman, and G. Violante (2020). The housing boom and bust: model meets evidence. *Journal of Political Economy*, Forthcoming.

Kaplan, G. and G. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica 82*(4), 1199–1239.

Lowry, S. (2014, August). Itemized tax deductions for individuals: data analysis. Crs report for congress, Congressional Research Service.

Piazzesi, M. and M. Schneider (2016). Housing and macroeconomics. *Handbook of Macroeconomics 2*, 1547–1640.

Sommer, K. and P. Sullivan (2018). Implications of US tax policy for house prices, rents, and homeownership. *American Economic Review 108*(2), 241–74.

Storesletten, K., C. I. Telmer, and A. Yaron (2004). Consumption and risk sharing over the life cycle. *Journal of Monetary Economics 51*(3), 609–633.

Tauchen, G. (1986). Finite state Markov-chain approximations to univariate and vector autoregressions. *Economics Letters 20*(2), 177–181.

The Tax Foundation (2013). U.S. federal individual income tax rates history, 1862-2013 (nominal and inflation-adjusted brackets).

Ventry, D. J. (2010). The accidental deduction: a history and critique of the tax subsidy for mortgage interest. *Law and Contemporary Problems 73*(1), 233–284.

1.A Equilibrium definitions

1.A.1 Stationary equilibria

Households are heterogeneous with respect to age $j \in \mathcal{J} \equiv \{1, 2, ..., J\}$, permanent earnings $z \in \mathcal{Z} \equiv \mathbb{R}_{++}$, mortgage $m \in \mathcal{M} \equiv \mathbb{R}_{+}$, owner-occupied housing $h \in \mathcal{H} \equiv \{0, \underline{\mathrm{h}}, ..., \bar{h} = \bar{s}\}$, and cash-on-hand $x \in \mathcal{X} \equiv \mathbb{R}_{++}$. Let $\mathcal{U} \equiv \mathcal{Z} \times \mathcal{M} \times \mathcal{H} \times \mathcal{X}$ be the non-deterministic state space with $\mathbf{u} \equiv (z, m, h, x)$ denoting the vector of individual states. Let $\mathbf{B}(\mathbb{R}_{++})$ and $\mathbf{B}(\mathbb{R}_{+})$ be the Borel σ-algebras on $\mathbb{R}_{++}$ and $\mathbb{R}_{+}$, respectively, and $P(\mathcal{H})$ the power set of $\mathcal{H}$, and define $\mathscr{B}(\mathcal{U}) \equiv \mathbf{B}(\mathbb{R}_{++}) \times \mathbf{B}(\mathbb{R}_{+}) \times P(\mathcal{H}) \times \mathbf{B}(\mathbb{R}_{++})$. Further, let $\mathbb{M}$ be the set of all finite measures over the measurable space $(\mathcal{U}, \mathscr{B}(\mathcal{U}))$. Then $\Phi_j(U) \in \mathbb{M}$ is a probability measure defined on subsets $U \in \mathscr{B}(\mathcal{U})$ that describes the distribution of individual states across agents of age $j \in \mathcal{J}$. Finally, denote the time-invariant fraction of the population of age $j \in \mathcal{J}$ by Π_j.

Stationary equilibrium with MID

Definition 1. A stationary recursive competitive equilibrium with MID ($\tau^m = 1$) is a collection of value functions $V_j(\mathbf{u})$ with associated policy functions $\{c_j(\mathbf{u}), s_j(\mathbf{u}), h'_j(\mathbf{u}), m'_j(\mathbf{u}),$
$b'_j(\mathbf{u})\}$ for all j; prices $(p_h = 1, p_r)$; a quantity of total housing stock $\bar{H}$; government's total tax revenue $\overline{TR}$; a quantity of total bequests left behind $\overline{BQ}$; and a distribution of agents' states Φ_j for all j such that:

1. Given prices $(p_h = 1, p_r)$, $V_j(\mathbf{u})$ solves the Bellman equation (1.8) with the corresponding set of policy functions $\{c_j(\mathbf{u}), s_j(\mathbf{u}), h'_j(\mathbf{u}), m'_j(\mathbf{u}), b'_j(\mathbf{u})\}$ for all j.

2. Given $p_h = 1$, the rental price per unit of housing service p_r is given by equation (1.9).

3. The quantity of the total housing stock is given by the total demand

for housing services[20]

$$\bar{H} = \sum_{\mathcal{J}} \Pi_j \int_U s_j(\mathbf{u}) d\Phi_j(U).$$

4. The government's net tax revenue $\overline{TR}$ is given by equation (1.12).

5. Total bequests $\overline{BQ}$ are given by equation (1.13).

6. The distribution of states Φ_j is given by the following law of motion for all $j < J$

$$\Phi_{j+1}(\mathcal{U}) = \int_U Q_j(\mathbf{u}, \mathcal{U}) d\Phi_j(U),$$

   where $Q_j : \mathcal{U} \times \mathscr{B}(\mathcal{U}) \rightarrow [0,1]$ is a transition function that defines the probability that a household at age j transits from its current state $\mathbf{u}$ to the set $\mathcal{U}$ at age $j+1$.

Stationary equilibrium without MID

Definition 2. A tax neutral stationary recursive competitive equilibrium without MID ($\tau^m = 0$) is a collection of value functions $V_j(\mathbf{u})$ with associated policy functions $\{c_j(\mathbf{u}), s_j(\mathbf{u}), h'_j(\mathbf{u}), m'_j(\mathbf{u}), b'_j(\mathbf{u})\}$ for all j; prices (p_h, p_r); a quantity of the total housing stock H; a parameter governing the average labor income tax level λ; a bequest parameter γ; and a distribution of agents' states Φ_j for all j such that:

1. Given prices (p_h, p_r) and parameters (γ, λ), $V_j(\mathbf{u})$ solves the Bellman equation (1.8) with the corresponding set of policy functions $\{c_j(\mathbf{u}), s_j(\mathbf{u}), h'_j(\mathbf{u}), m'_j(\mathbf{u}), b'_j(\mathbf{u})\}$ for all j.

2. Given p_h, the rental price per unit of housing service p_r is given by equation (1.9).

[20] We assume a perfectly elastic supply of both owner-occupied housing and rental units in the initial steady state. This implies that supply always equals demand, and we thus have market clearing.

3. The housing market clears:

$$H = \bar{H}$$
$$\text{where} \quad H = \sum_{\mathcal{J}} \Pi_j \int_U s_j(\mathbf{u}) d\Phi_j(U)$$
$$\text{and} \quad \bar{H} \text{ is the housing stock from the equilibrium with MID.}$$

4. The government's net tax revenue is the same as in the steady state with MID:

$$TR = \overline{TR}$$
$$\text{where} \quad TR \text{ is given by equation (1.12)}$$
$$\text{and} \quad \overline{TR} \text{ is the tax revenue from the equilibrium with MID.}$$

5. The bequest parameter γ is the solution to

$$BQ - \overline{BQ} = \sum_{j=1}^{J-1} \Pi_j \phi_j \int_U \gamma z(\mathbf{u}) d\Phi_j(U)$$
$$\text{where} \quad BQ \text{ are given by equation (1.13)}$$
$$\text{and} \quad \overline{BQ} \text{ are the bequests from the equilibrium with MID.}$$

6. Distributions of states Φ_j are given by the following law of motion for all $j < J$

$$\Phi_{j+1}(\mathcal{U}) = \int_U Q_j(\mathbf{u}, \mathcal{U}) d\Phi_j(U),$$

1.A.2 Transitional equilibrium

Let $\Phi_{tr,jt}(U_t) \in \mathbb{M}$ be a probability measure defined on subsets $U_t \in \mathscr{B}(\mathcal{U})$ that describes the distribution of individual states across agents of age $j \in \mathcal{J}$ at time period t.

Definition 3. Given a sequence of mortgage interest deductibility parameters $\{\tau_t^m\}_{t=1}^{t=\infty}$ and initial conditions $\Phi_{tr,j1}(U_1)$ for all j, a tax neutral transitional recursive competitive equilibrium is a sequence of value functions $\{V_{jt}(\mathbf{u})\}_{t=1}^{t=\infty}$ with associated policy functions $\{c_{jt}(\mathbf{u}), s_{jt}(\mathbf{u}), h'_{jt}(\mathbf{u}), m'_{jt}(\mathbf{u}), b'_{jt}(\mathbf{u})\}_{t=1}^{t=\infty}$ for all j; a sequence of prices

$\{(p_{h,t}, p_{r,t})\}_{t=1}^{t=\infty}$; a sequence of quantities of total housing demand $\{H_t\}_{t=1}^{t=\infty}$; a sequence of parameters governing the average labor income tax level $\{\lambda_t\}_{t=1}^{t=\infty}$; a sequence of bequest parameters $\{\gamma_t\}_{t=1}^{t=\infty}$; and a sequence of distributions of agents' states $\{\Phi_{tr,jt}\}_{t=1}^{t=\infty}$ for all j such that:

1. Given prices $(p_{h,t}, p_{r,t})$ and parameters (γ_t, λ_t), $V_{jt}(\mathbf{u})$ solves the Bellman equation with the corresponding set of policy functions $\{c_{jt}(\mathbf{u}), s_{jt}(\mathbf{u}), h'_{jt}(\mathbf{u}), m'_{jt}(\mathbf{u}), b'_{jt}(\mathbf{u})\}$ for all j and t.

2. Given $p_{h,t}$ and $p_{h,t+1}$, the rental price per unit of housing service is $p_{r,t}$ for all t.

3. The housing market clears:

$$\begin{aligned} & H_t = \bar{H} \quad \forall t \\ \text{where} \quad & H_t = \sum_{\mathcal{J}} \Pi_j \int_{U_t} s_{jt}(\mathbf{u}) d\Phi_{tr,jt}(U_t) \quad \forall t \\ \text{and} \quad & \bar{H} \text{ is the housing stock from the equilibrium with MID.} \end{aligned}$$

4. The government's net tax revenue is the same as in the steady state with MID:

$$\begin{aligned} & TR_t = \overline{TR} \quad \forall t \\ \text{where} \quad & TR_t \ \text{ is the total tax revenue in period } t, \forall t \\ \text{and} \quad & \overline{TR} \ \text{ is the tax revenue from the equilibrium with MID.} \end{aligned}$$

5. The bequest parameter γ_t is the solution to:

$$BQ_t - \overline{BQ} = \sum_{j=1}^{J-1} \Pi_j \phi_j \int_{U_t} \gamma_t z(\mathbf{u}) d\Phi_{tr,jt}(U_t) \quad \forall t$$

$$\begin{aligned} \text{where} \quad & BQ_t \ \text{ is the value of bequests in period } t, \forall t \\ \text{and} \quad & \overline{BQ} \ \text{ are the bequests from the equilibrium with MID.} \end{aligned}$$

6. Distributions of states $\Phi_{tr,jt}$ are given by the following law of motion for all $j < J$ and t:

$$\Phi_{tr,j+1,t+1}(\mathcal{U}) = \int_{U_t} Q_{tr,jt}(\mathbf{u}, \mathcal{U}) d\Phi_{tr,jt}(U_t),$$

where $Q_{tr,jt} : \mathcal{U} \times \mathscr{B}(\mathcal{U}) \to [0,1]$ is a transition function that defines the probability that a household of age j at time t transits from its current state $\mathbf{u}$ to the set $\mathcal{U}$ at age $j+1$ and time $t+1$.

1.B Computational method

We discretize the state space by choosing a finite grid for permanent earnings $Z_j \equiv \{z_{j,1}, ..., z_{j,N_Z}\}$ and cash-on-hand $X \equiv \{x_1, ..., x_{N_X}\}$.[21] Permanent earnings are age specific with $N_Z = 9$ grid points. We set the number of cash-on-hand grid points N_X to 30. Moreover, we take into account the concavity of the value function by letting the spacing between two grid points increase with the level of cash-on-hand. Housing is assumed to be available in discrete sizes only and we let the grid for housing be $H \equiv \{0, h_1, ..., h_{N_H}\}$ where h_1 is calibrated and $N_H = 9$.

To solve for the value and policy functions, we use the general generalization of the endogenous grid method G^{2}EGM by Druedahl and Jørgensen (2017). The method allows for occasionally binding constraints and non-convexities, while reaping the speed benefits associated with the traditional EGM as in Carroll (2006).

We approximate expectations to solve for the value and policy functions. The transitory earnings shocks are approximated by five Gauss-Hermite quadrature nodes, whereas the permanent earnings shocks are approximated using Markov chains. We use the method in Tauchen (1986), but allow the support for shocks to fan out over the life cycle (see, e.g., Storesletten et al. (2004)). For each age, we let the outermost grid points be $m_Z = 3$ standard deviations from the mean. For simulation purposes, we draw both shocks from their respective continuous distributions. To avoid extrapolation of permanent shocks outside the $m_Z = 3$ standard deviation bound, we force permanent income to be on the outermost grid point whenever necessary.

Similar to the traditional EGM, we use grids for the post-decision states to solve for the value and policy functions. The post-decision states in our model are bonds $b' \in \mathbb{R}_+$, mortgages $m' \in \mathcal{M} \equiv \mathbb{R}_+$, and housing $h' \in \mathcal{H}$. We force m' to be on grid whenever the household chooses a

[21] We do, however, allow households to have permanent earnings z and cash-on-hand x off grid. We linearly interpolate in cases where z and x are off grid.

positive amount of bonds, and mortgages are not given by a constraint. For computational convenience, we let b'_y and ltv' be post-decision states instead of b' and m', respectively, where b'_y denotes bonds as a fraction of earnings and ltv' denotes loan-to-value.[22]

Let ϵ be a very small positive number. We choose a finite grid for bonds over earnings $B_y \equiv \{b_{y,1} = 0, b_{y,2} = \epsilon, b_{y,3}, ..., b_{y,N_B}\}$ where $N_B = 25$ and the grid points are denser at lower levels of bonds over earnings. The finite grid for loan-to-value is $LTV \equiv \{ltv_1 = 0, ltv_2 = \epsilon, ..., (1-\theta-\epsilon), (1-\theta), (1-\theta+\epsilon), ..., ltv_{N_{LTV}}\}$ where $N_{LTV} = 21$ and θ is the down-payment requirement. Between ltv_2 and $(1-\theta-\epsilon)$ spacing is linear. Spacing is also linear between $(1-\theta+\epsilon)$ and $ltv_{N_{LTV}}$. We allow policy functions for b'_y and ltv' to be off grid by using linear interpolation.

From the definition of the finite grid LTV, we can see how the alternative formulation of post-decision states is particularly convenient in the case of mortgages. First, we ensure that the loan-to-value requirement is on the discretized grid. Second, we can easily specify loan-to-value levels that are very close to the occasionally binding constraints. Both these features help facilitate more efficient and accurate solutions.

To solve for the equilibrium, we simulate $150,000$ households for J periods. When aggregating, each age group is assigned a weight Π_j, where the weight reflects the true population density in the U.S. Households are born with some initial assets. During their lives, they receive earnings shocks from continuous distributions, along with some bequests, at the beginning of each period. Households then pay taxes before they make their choices.

All policy reforms are unexpected and we adjust individual states for changes in the house price and taxes. Specifically, cash-on-hand x needs to be adjusted due to the fact that (i) the value of the house falls; (ii) the property tax payment falls; (iii) there are lower tax deductions due to changes in the MID and lower property taxes; (iv) there are changes in the tax level parameter λ; (v) there are changes in the bequest parameter γ; and (vi) there are losses incurred by rental companies. In addition, we need to adjust for changes in the loan-to-value due to a lower house price.

At any time t during the transition, new households enter the economy

[22] Note that both b' and m' can easily be backed-out from b'_y and ltv', for given earnings y, housing h', and house price p_h.

and replace the households that die between periods $t-1$ and t. We assume that newborns are hit by the same sequences of exogenous earnings shocks as the households they replace.

1.C Solution algorithm

1.C.1 Steady state

Solving the initial steady state with MID ($\tau^m = 1$):

1. Impose house price $p_h = 1$ and compute p_r from equation (1.9).
2. Solve the household problem recursively, and obtain the value and policy functions.
3. Simulate using optimal decision rules.
4. Use simulated values to compute the total housing stock $\bar{H}$, the government's total tax revenue $\overline{TR}$, and total bequests $\overline{BQ}$. From the simulation, we also get the distribution of agents' states Φ_j for all j.

Let λ_{init} be the parameter value of the labor income tax level in the initial steady state. Then, solving the new tax and bequest neutral steady state without MID ($\tau^m = 0$) can be divided into 2 stages. In the first stage, we solve the steady state without MID given that $\lambda = \lambda_{init}$ and $\gamma = 0$, i.e., we do not impose tax neutrality and do not consider changes in the amount of bequest:

1. Guess p_h and compute p_r.
2. Recursively solve for the value and policy functions, and simulate using optimal decision rules. Use simulated values to compute the total housing demand H.
3. Compute excess demand for housing $ED_H = H - \bar{H}$.
   (a) If $|ED_H|$ is larger than some tolerance level, update p_h using bisection and return to step 1.

(b) If $|ED_H|$ is within the tolerance level, convergence in the first stage is achieved. Denote the equilibrium house price under stage 1 as $\hat{p}_h$.

In the second stage, we solve for the tax and bequest neutral steady state:

1. Guess (p_h, λ, γ), where the first guess is $p_h = \hat{p}_h + \epsilon_{p_h}$, $\lambda = \lambda_{init} + \epsilon_\lambda$, and $\gamma = 0 + \epsilon_\gamma$.

2. Given p_h, compute p_r.

3. Recursively solve for the value and policy functions, and simulate using optimal decision rules. Use simulated values to compute the total housing demand H, government's total tax revenues TR, total bequests distributed $\widehat{BQ}$, and total bequests collected BQ.

4. Compute excess demand for housing, excess government tax revenue, and the excess bequest, ED_H, $ED_{TR} = TR - \overline{TR}$, and $ED_{BQ} = (BQ - \overline{BQ}) - \widehat{BQ}$, respectively.

   (a) If $|ED_H|$, $|ED_{TR}|$, and/or $|ED_{BQ}|$ are larger than some tolerance levels, update the guess for (p_h, λ, γ) using the rule $q' = q + ED_k * \epsilon_q$ where $q \in \{p_h, \lambda, \gamma\}$ and $k = H$ if $q = p_h$, $k = TR$ if $q = \lambda$ and $k = BQ$ if $q = \gamma$. Return to step 2.

   (b) If all of $|ED_H|$, $|ED_{TR}|$, and $|ED_\gamma|$ are within the tolerance levels, convergence is achieved.

1.C.2 Transition

Let $\Phi_{init,j}$ be the distribution of households' states in the initial steady state, and let λ_{new} and γ_{new} be the equilibrium λ and γ from the tax and bequest neutral steady state without MID. Further, let t denote the transition period, and assume that the economy is in the new steady state in $t = T + 1$. Choose T large enough so that by increasing T the transition path is unaltered.[23] The solution algorithm for the transitional equilibrium can be described in two stages. In the first stage, we solve for the transitional equilibrium assuming $\lambda_t = \lambda_{new}$ and $\gamma_t = \gamma_{new}$ $\forall t \in \mathcal{T} \equiv \{1, ..., T\}$:

[23]We set $T = J + 5$.

1. Guess $\{p_{h,t}\}_{t=1}^{t=T}$ and compute $\{p_{r,t}\}_{t=1}^{t=T}$.

2. Recursively solve for the value and policy functions for all ages $j \in \mathcal{J}$ and time periods $t \in \mathcal{T}$. To solve for value and policy functions at time period $t = T$, assume that the value and policy functions in $t = T + 1$ are those from the new steady state with neutrality.

3. Given the price $p_{h,1}$ and parameters γ_1 and λ_1, for each $j \in \mathcal{J}$, adjust the initial individual states such that the initial distribution $\Phi_{init,j}$ reflects unexpected changes in the house price, the tax level, and bequests from the initial steady state.

4. Simulate using the adjusted initial distribution and optimal decision rules. Use simulated values to compute the sequence of total housing demand $\{H\}_{t=1}^{t=T}$.

5. Compute the sequence of excess demand for housing $\{ED_{H,t}\}_{t=1}^{t=T}$, and the Euclidean norm of this sequence.

    (a) If the norm is larger than some tolerance level, update $\{p_{h,t}\}_{t=1}^{t=T}$ using the rule $p'_{h,t} = p_{h,t} + ED_{H_t} * \epsilon_{p_h}$ for all $t \in \mathcal{T}$ and return to step 1.

    (b) If the norm is within the tolerance level, convergence in the first stage is achieved. Denote the equilibrium house prices under stage 1 $\hat{p}_{h,t}$ for all $t \in \mathcal{T}$.

In the second stage, we solve for the tax neutral transitional equilibrium:

1. Guess $\{(p_{h,t}, \lambda_t, \gamma_t)\}_{t=1}^{t=T}$, where the first guess is $p_{h,t} = \hat{p}_{h,t}$, $\lambda_t = \lambda_{new}$, and $\gamma_t = \gamma_{new}$ for all $t \in \mathcal{T}$.

2. Given $\{p_{h,t}\}_{t=1}^{t=T}$, compute $\{p_{r,t}\}_{t=1}^{t=T}$.

3. Recursively solve for the value and policy functions for all ages and time periods, adjust the initial individual states such that the initial distribution $\Phi_{init,j}$ reflects unexpected changes in the house price, the tax level and bequests from the initial steady state, and simulate using the adjusted initial distribution and optimal decision rules. Use simulated values to compute the sequences of total housing demand $\{H\}_{t=1}^{t=T}$, the government's total tax revenues $\{TR\}_{t=1}^{t=T}$, the

total bequests distributed $\{\widehat{BQ}\}_{t=1}^{t=T}$, and the total bequests collected $\{BQ\}_{t=1}^{t=T}$.

4. Compute the sequences of excess demand for housing, excess government tax revenue, and excess bequests, $\{ED_{H,t}\}_{t=1}^{t=T}$, $\{ED_{TR,t}\}_{t=1}^{t=T}$, and $\{ED_{BQ,t}\}_{t=1}^{t=T}$, respectively. Compute the Euclidean norm of all three sequences.

   (a) If the norm of either sequence is larger than some tolerance level, update the guess $\{(p_{h,t}, \lambda_t, \gamma_t)\}_{t=1}^{t=T}$ using the rule $q' = q + ED_k * \epsilon_q$ for all $t \in \mathcal{T}$, where $q \in \{p_{h,t}, \lambda_t, \gamma_t\}$ and $k = H_t$ if $q = p_{h,t}$, $k = TR_t$ if $q = \lambda_t$, and $k = BQ_t$ if $q = \gamma_t$. Return to step 2.

   (b) If all norms are within the tolerance levels, convergence is achieved.

1.D Labor income process

1.D.1 Data sample

Equations (1.18) and (1.19) are estimated using PSID data for the survey years 1970 to 1992. Following Cocco et al. (2005), we drop households where the head was i) a nonrespondent, ii) part of the Survey of Economic Opportunities subsample, iii) disabled or retired, iv) a student, or v) a housewife. Due to few female headed households, we exclusively focus on households with male heads.

In line with Guvenen (2009), we further restrict the sample by only keeping households for which i) earnings are strictly positive, ii) annual hours worked by head are between 520 (10 hours per week) and 5110 (14 hours a day, everyday), iii) the head's average hourly wage is between \$2 and \$400 (inclusive) in 1993 dollars, where we adjust the bounds backwards using the growth rate in average weekly earnings from "Current Employment Statistics" published by the Bureau of Labor Statistics. Series ID: CES0500000030. iv) the head is between 20 and 64 years old, and v) the head appears in the sample in at least 15 out of 23 possible survey years.

1.D.2 Estimation

In order to simulate the exogenous earnings process according to equations (1.18) and (1.19), we estimate the deterministic earnings profile $g(j)$ and the variances of the fixed-effect component σ_α^2, the permanent shock σ_η^2, and the transitory shock σ_ν^2. Estimating the deterministic wage component $g(j)$ is done in two steps. First, we estimate it on an annual basis, and then we convert it to suit the model period length of three years.

Step 1: Using the yearly observations in the data, we estimate a yearly version of the deterministic component. That is, we estimate $g_{annual}(age)$, where $age \in \{20, 21, ..., 64\}$. We regress $\log(y_i)$ on dummies for age (not including the youngest age), marital status, family composition (number of family members besides head and, potentially, wife), and a dummy for whether the agent has a college education or not. We control for household fixed effects by running a linear fixed effect regression. We fit a third-order polynomial to the predicted values of this regression, which gives us the estimate of the *annual* deterministic earnings profile $\hat{g}_{annual}(age)$.

Step 2: We convert annual estimates to three-year periods as follows

$$\hat{g}(j) = \hat{g}_{annual}(j * 3 + 21) \quad \text{for } j \in [1, J_{ret}]. \tag{1.21}$$

Equation (1.21) states that the deterministic earnings in period $j = 1$ are the annual deterministic earnings at adult age 24 and the earnings in period $j = J_{ret}$ are the annual earnings at adult age 63. As such, the deterministic earnings in period j are equal to those of the middle adult age that period j is assumed to represent.

With an estimate of the deterministic earnings profile at hand, the variances of the transitory (σ_ν^2) and permanent (σ_η^2) shocks are estimated in a similar fashion as in Carroll and Samwick (1997). Define $\log(y_{ij}^*)$ as the logarithm of earnings less the household fixed component and the deterministic earnings path

$$\begin{aligned}\log(y_{ij}^*) &\equiv \log(y_{ij}) - \hat{\alpha}_i - \tilde{g}(j) \\ &= n_{ij} + \nu_{ij} \qquad\qquad \text{for } j \in [1, J_{ret}],\end{aligned}$$

where the equality follows from equation (1.18). Since we have three-year periods in the model, we define $\log(y_{ij})$ as the sum of earnings from the three adult ages to which the model period corresponds. For example, $\log(y_{i1}) = \log(\sum_{age=23}^{25} y_{i,age}^{annual})$. Similarly, $\tilde{g}(j)$ is defined as the sum of the annual deterministic earnings components, for example $\tilde{g}(1) = \log\left(\sum_{age=23}^{25}\exp(\hat{g}_{annual}(age))\right)$. Next, define household i's d-period difference in $\log(y^*_{ij})$ as

$$\begin{aligned} r_{id} &\equiv \log(y^*_{i,j+d}) - \log(y^*_{ij}) \\ &= n_{i,j+d} + \nu_{i,j+d} - n_{ij} - \nu_{i,j} \\ &= n_{i,j+1} + n_{i,j+2} + ... + n_{i,j+d} + \nu_{i,j+d} - \nu_{i,j}. \end{aligned}$$

In the last step, we recursively apply equation (1.19). Using that the transitory and permanent shocks are i.i.d., it follows that

$$\begin{aligned} \text{Var}(r_{id}) &= \text{Var}(n_{i,j+1}) + \text{Var}(n_{i,j+2}) + ... + \text{Var}(n_{i,j+d}) \\ &\quad + \text{Var}(\nu_{i,j+d}) + \text{Var}(\nu_{i,j}) \\ &= d\,\sigma^2_\eta + 2\,\sigma^2_\nu. \end{aligned}$$

We estimate these variances by running an OLS regression of $\text{Var}(r_{id}) = r^2_{id}$ on d and a constant term. Then, the coefficient of d is our estimate of the variance of the permanent shock, whereas the constant term divided by two is our estimate of the variance of the transitory shock.

Finally, the estimate of σ^2_α is the residual variance in period $j = 1$ as follows

$$\hat{\sigma}^2_\alpha = \text{Var}\left(\log(y_{i1}) - \tilde{g}(1)\right) - \hat{\sigma}^2_\eta - \hat{\sigma}^2_\nu.$$

1.D.3 Variable definitions

Age of head is constructed by taking the first observed age and then adding the number of years between a given survey year and the first survey in which the individual was observed. This is to avoid non-changes and two-year jumps in the age variable between two consecutive survey years. The variable name of age is V20651 in the 1992 PSID survey.

CPI is taken from the BLS. We use the historical CPI for all urban consumers (CPI-U), U.S. city average, all items.

Family composition is the number of family members besides head and, potentially, wife. We define it as family size less adults. Family size is the number of members in the family unit at the time of an interview. Adults are defined as the number of major adults (head and wife only). The variable names are V20398 and V20397 in the 1992 PSID survey for family size and adults, respectively.

Head's education is divided into two groups: households with a college degree and households with no college degree. Between 1970 to 1990, we define the education groups by using the categorical groups defined in the PSID. For example, in the 1990 survey we use the variable name V18898, and define that no college consists of levels 1 to 6, and college comprises levels 7 and 8. After 1990, we use a variable for years of completed education (variable name V21504 in 1992 survey). Then, no college households comprise levels 0 to 15 and households with a college degree comprise levels 16 and 17. We drop observations where individuals have no appropriate answer (NA or don't know) and individuals who before the 1984 survey answered "Could not read or write; DK grade and could not read or write".

Head's annual labor hours are the total annual work hours on all jobs including overtime. The variable name is V20344 in the 1992 PSID survey.

Head's average hourly wage is computed as the head's wage divided by the head's annual labor hours.

Household earnings y_{ij} are the sum of labor income for both head and wife. Earnings are deflated with the CPI using 1992 as the base year. Labor income is defined as the sum of salary income, bonuses, overtime, commissions, the labor part of farm, business, market gardening, roomers and boarders income, and income from professional practice or trade. The variable names are V21484 and V20436 in the 1992 PSID survey for head and wife, respectively.

The maximum allowable benefit during retirement, B^{max} in equation (1.20), is computed using data from the Social Security Administration (SSA). Specifically, we use the maximum monthly benefit level that was available for a person retiring at age 66 in 1992 ($1,113) and multiply it by twelve to get a yearly benefit level. We adjust the yearly level for the difference in the SSA's average wage per worker in 1992 ($22,002) and the average earnings in the model.

1.E A grandfather policy

To investigate the effects of a removal policy in which we discriminate between cohorts, we study the effects of a policy where new households are not allowed to deduct mortgage interest payments, while existing households can continue to do so. We refer to this policy as the *grandfather* policy. Figure 1.8 shows the transition paths for the house price, the rental price, the average marginal tax rate, and the bequest parameter.

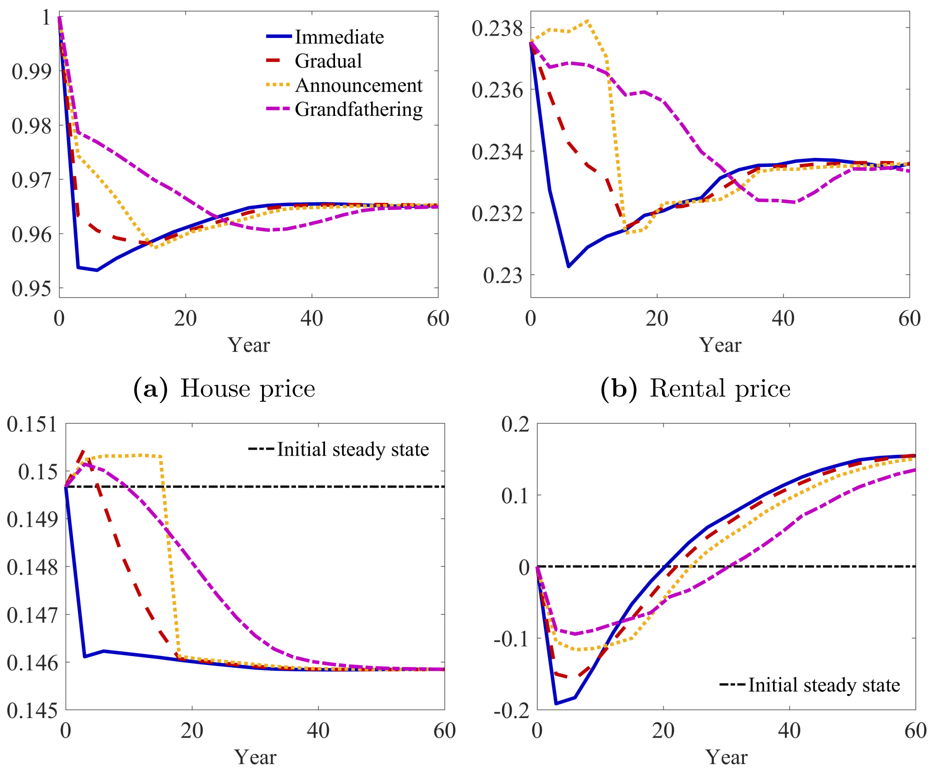

(a) House price

(b) Rental price

(c) Average marginal labor income tax rate before deductions

(d) Bequest rate γ (%)

Figure 1.8: Short-run dynamics from removing the MID, across policies
Note: All policies are implemented unexpectedly and households have perfect foresight of the transition paths. The respective panels show how the house price, the rental price, the average marginal tax rate before deductions, and the bequest rate behave in the short run, in response to the changes in the deductibility rate. The rental price corresponds to a three-year (one model period) cost of renting.

Naturally, the convergence for the grandfather policy is slower than for the alternative policies. There is also a smaller immediate fall in the house price, as only the households that enter the economy are directly affected by the MID removal. The slower fall in the house price leads to a correspondingly slower fall in the rental price. Under the grandfather policy the labor income tax rate increases initially, as the government still spends large amounts on interest deductions and the revenue from property taxes falls. As new cohorts replace older cohorts, the labor income tax level slowly declines towards the lower level of the new steady state. The value of bequests falls immediately under this policy as well, since the house price decreases. Over time, this amount slowly converges to the new steady state.

Table 1.10 presents the average CEV, and the fraction in favor, for the four policies. The grandfather policy is able to limit the welfare losses quite substantially for many homeowners, which leads to an average welfare effect close to that of the immediate policy. However, the fraction of households with a welfare gain is still low. The reason for this low support is that a significant share of renters are not in favor of the reform.

	Immediate	Gradual	Announcement	Grandfather
Mean CEV (%)	-0.40	-0.52	-0.52	-0.38
Fraction in favor	0.39	0.35	0.27	0.31

Table 1.10: Welfare effects for households alive in the first period of the transition
Note: The fraction in favor is the fraction of households with a CEV greater than or equal to zero. For a description of CEV (%), see *Note* below Figure 1.5.

Figure 1.9 displays the distribution of CEV for the four policies. Compared to the other policies, the grandfather policy has a higher house price, and a relatively high rental price and taxes for most of the transition. All these effects contribute to the lower welfare of renters, and combined with the initial drop in bequests, pushing some of these households into negative CEV territory. Similar to the other policies, most homeowners experience welfare losses from the grandfather reform. Homeowners are negatively affected by the fall in the house price and the instantaneous increase in the labor income tax level and fall in the bequest rate. However, since they can still deduct mortgage interest payments, their welfare losses are limited, especially for households with high earnings and high LTV-ratios.

Overall, the analysis of the welfare effects of the grandfather policy is similar to that of other more gradual policies. By removing the MID slowly, the welfare distribution is compressed. The households who lose the most from a repeal of the MID realize smaller welfare losses, and the households who benefit the most experience smaller welfare gains.

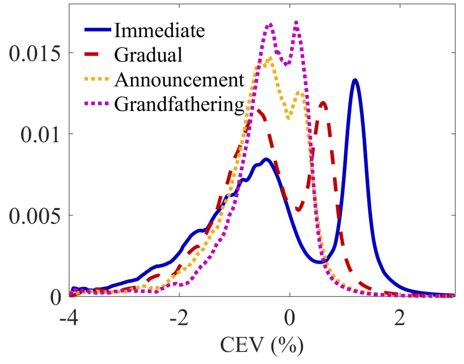

Figure 1.9: Distributions of short-run welfare effects across policies, including grandfathering

Note: Distributions of welfare effects for all policies, for households alive in the first period of the transition. For a description of CEV (%) see *Note* below Figure 1.5.

1.F Welfare effects: equilibrium assumptions

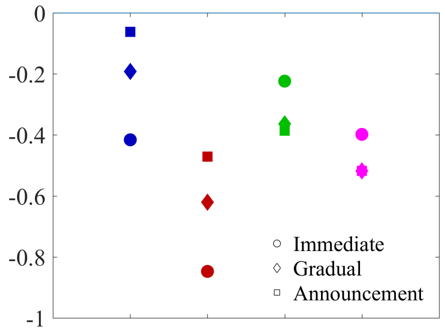

Rental and house prices adjust	-	✓	✓	✓
Tax neutrality	-	-	✓	✓
Bequests adjust	-	-	-	✓

Figure 1.10: Short-run welfare effects under different equilibrium assumptions
Note: The first column shows the mean CEV for those alive in the first period of the transition, when we only consider the direct effect of removing the MID. We account for rental companies' losses in the first period of the transition when we allow for prices to change. For a description of CEV (%) see *Note* below Figure 1.5.

1.G Tax Cuts and Jobs Act: long-run results

	Baseline			Tax Cuts and Jobs Act			Tax Cuts and Jobs Act, no cap		
	MID	No MID	Difference (%)	MID	No MID	Difference (%)	MID	No MID	Difference (%)
House price	1	0.965	-3.47	0.988	0.969	-1.85	0.997	0.975	-2.13
Rental price	0.238	0.234	-1.66	0.236	0.234	-0.88	0.237	0.235	-1.02
Homeownership rate	0.70	0.71	1.88	0.67	0.71	5.71	0.65	0.70	6.96
Fraction ever-owner	0.88	0.89	1.59	0.88	0.89	1.60	0.87	0.88	2.02
Mean owned house size	215	211	-2.15	215	211	-1.67	217	212	-2.46
Mean LTV	0.36	0.31	-12.09	0.31	0.31	0.73	0.30	0.31	1.15
Mean mortgage	74	60	-19.29	63	59	-5.43	64	59	-6.83
Mean bond holdings	20.6	21	1.81	20.9	21.3	1.85	20.6	21.3	3.40
Mean marginal tax rate	0.150	0.146	-2.59	0.150	0.149	-0.16	0.150	0.149	-0.65
Mean bequest collected	152	158	3.57	156	159	1.77	156	159	2.15
Itemization rate	0.53	0.19	-64.50	0.09	0	-100	0.12	0.02	-79.99

Table 1.11: Long-run effects on prices and aggregates of removing the MID, baseline versus Tax Cuts and Jobs Act

Note: The table shows steady-state results based on three different initial tax systems. The first tax system called “Baseline” simply reiterates the results from Table 1.4. In the “Tax Cuts and Jobs Act”, we multiply the baseline standard deduction by 1.9, and the maximum deductions for the sum of state and local income taxes and property taxes are set to 10,000 in 2018 dollars. In the last tax system, we multiply the baseline standard deduction by 1.9, but there is no change in the cap. The first column within each of these tax systems shows prices and aggregate measures in the initial steady state with MID, whereas the second column shows the corresponding values in the steady state without MID. The rental price corresponds to a three-year (one model period) cost of renting. “Fraction ever-owner” is the fraction of households that own a house at some point during their life. The mean house size, LTV, and the mortgage level are conditional on owning. The mean owned house size, mortgage, bond holdings, and bequest collected are in 1000’s of 2013 dollars. The mean marginal tax rate is gross of deductions.

Chapter 2

Monetary policy and the mortgage market*

*I am grateful for advice from and discussions with Per Krusell and Kurt Mitman. I also wish to thank Adrien Auclert, Timo Boppart, Tobias Broer, Mitch Downey, Richard Foltyn, John Hassler, Markus Karlman, Kasper Kragh-Sørensen, John Kramer, Kieran Larkin, Claire Thürwächter, Roine Vestman, Magnus Åhl, and seminar participants at BI Norwegian Business School, Banco de Portugal, Bank of Canada, Royal Holloway, University of London, Sveriges Riksbank, the European Central Bank, the Federal Reserve Bank of New York, the Federal Reserve Bank of San Francisco, the IIES, the Norwegian Central Bank, the Norwegian School of Economics (NHH), University of Alicante, University of Bristol, University of Guelph, University of Wisconsin–Madison, and Uppsala University. I gratefully acknowledge funding from Handelsbanken's Research Foundations and Torsten Söderbergs Stiftelse. Computations were performed on resources provided by SNIC through Uppsala Multidisciplinary Center for Advanced Computational Science (UPPMAX) under Project SNIC 2019/8-311. All errors are my own.

2.1 Introduction

Monetary policy affects many households through its effect on mortgage interest rates and house prices. During the Great Recession, both mortgage interest rates and house prices were focal points for many central banks. In the U.S. over 40 percent of households have a mortgage and the outstanding mortgage debt surpassed USD 15.5 trillion in 2019, which corresponds to about 70 percent of GDP. Further, housing is by far the largest asset on most American households' balance sheets. In this paper, I explore how changes in mortgage interest rates affect aggregate demand, and what role house prices play in the transmission of monetary policy.

Recent empirical studies indicate that there are two important components related to mortgages and housing that influence how strongly aggregate demand responds to changes in the interest rate.[1] First, the extent to which households are exposed to changes in mortgage interest rates matters, which depends both on the type of mortgage contracts that are used and the behavioral response by homeowners who choose to refinance their mortgage. Second, the endogenous response in house prices depends on the structure of the mortgage market.

I quantify the role of changes in mortgage interest rates and house prices in the transmission of monetary policy. For this purpose, I build a heterogeneous-agent life-cycle model where the mortgage and housing markets are modeled in detail, and where house prices are endogenous. I further explore how the aggregate response in consumption depends on the type of mortgage contract that is used. Specifically, I compare responses when the available mortgage is a fixed-rate 30-year mortgage, which is the most commonly used type in the U.S., to a setting where mortgages have adjustable rates.

Following an expansionary monetary policy shock, I find that the interaction of changes in mortgage interest rates and house prices amplifies the response in aggregate consumption. The amplification is mainly driven by constrained homeowners who refinance their mortgages to smooth consumption. When mortgages have adjustable rates, house prices increase substantially more in response to a reduction in the interest rate, relative to when mortgages have fixed rates. The stronger response in

[1] See, for example, Calza et al. (2013), Cloyne et al. (2019), Di Maggio et al. (2017), and Flodén et al. (2019).

house prices enables larger cash-outs among those who refinance and more households find refinancing optimal. As a result, the aggregate response in consumption is over six times as large under adjustable-rate mortgages as compared to when fixed-rate mortgages are used.

In the model, households choose how much to consume, whether to rent or own a house, their house size, mortgage financing, and savings in risk-free liquid bonds. Importantly, owned housing is illiquid, and markets are incomplete as households cannot insure against idiosyncratic earnings risks. There are two features of the housing market that create the illiquid nature of housing equity. First, households pay transaction costs to buy or sell a house. Second, if a homeowner wants to access its housing equity, it incurs refinancing costs. Additional features of the housing market include down-payment and payment-to-income requirements that have to be fulfilled when purchasing a home or refinancing a mortgage.

Since households cannot perfectly insure against earnings risks, there are households who are constrained due to poor earnings realizations. Further, since housing wealth is illiquid, there are also some relatively wealthy households that are constrained in their spending. These "wealthy hand-to-mouth" households can play an important role for the response of aggregate demand to shocks, as emphasized by Kaplan and Violante (2014) and Kaplan et al. (2018). In addition, the existence of mortgage financing allows for many households to be both relatively poor, i.e., have large mortgage balances, and have high exposures to changes in the interest rate. These indebted households tend to be young homeowners who recently bought their first home. As young households expect higher earnings in the future, they mainly save for precautionary reasons and have high marginal propensities to consume.

The calibrated model successfully matches the life-cycle profiles of homeownership, loan-to-value (LTV), housing wealth-to-earnings, and mortgage-to-earnings in the U.S. economy. The model also produces realistic distributions of liquid asset-to-earnings and net-worth-to-earnings in the lower ranges of the distributions, among the financially constrained households. A large share of households with low liquid asset-to-earnings ratios are relatively wealthy homeowners with most of their wealth invested in housing. In the model, about 32 percent of the wealthy households are liquidity constrained, while the share of hand-to-mouth households among the wealthy in the data is approximately 31 percent.[2] Thus, by including housing as an illiquid asset, with frictions measured in the data, the model is able to match the prevalence of wealthy, liquidity-constrained households. The model also successfully replicates the negative correlation between LTVs and liquid assets that is found in the data.

In order to study how households respond to a monetary policy shock, I use an empirically estimated path of the real interest rate from a shock of -100 basis points (bp) to the nominal interest rate. The estimated path of the real interest rate is the impulse response function from the identified Romer and Romer (2004) monetary policy shock in Auclert et al. (2020). The negative shock of 100bp to the nominal interest rate translates into an immediate reduction of the real interest rate of approximately 80bp. I assume perfect foresight of the path of the real interest rate following the shock, and the long-term mortgage interest rate is given by the expected future real interest rates and an exogenous credit spread.[3]

I find that when allowing for changes in mortgage interest rates and house prices, the aggregate response of consumption is amplified. I also find that the structure of the mortgage market impacts the effectiveness of monetary policy substantially. In the baseline scenario where households use 30-year mortgages with fixed rates, the aggregate immediate response of consumption to the expansionary real interest rate shock is 50 percent higher (0.09 percent vs 0.06 percent) due to changes in mortgage interest rates and house prices. This increase is primarily driven by homeowners, in particular by constrained households who refinance their mortgage to smooth consumption. Interestingly, I find that neither changes in house prices nor mortgage interest rates alone contribute significantly to the amplification, but it is the interaction that matters. This complementarity effect is due to an endogenous increase in house prices, which relaxes the LTV-constraint among refinancers and enables larger cash-outs. The reduction in the mortgage rate incentivizes the refinancing among constrained homeowners, and the increase in house prices further improves their consumption smoothing.

The increase in consumption due to the reduction in the real interest rate is accompanied by a decline in savings. Both the classical channel

[2]I assume that a household is liquidity constrained if it has less than half of monthly earnings in liquid savings.

[3]I assume that risk premia and thus credit spreads and term premia are constant.

of monetary policy, i.e., the intertemporal substitution channel, and the increase in consumption smoothing among constrained homeowners contribute to the intertemporal shift in consumption. The reduction in savings is characterized by a shift from liquid bonds into mortgages, among homeowners. When mortgages have fixed, long-term rates, the mortgage interest rate responds less than one-for-one with changes in the short-term policy rate. Further, only households who take up a new mortgage are affected by the lower mortgage interest rate. As a result, a temporary reduction in the risk-free rate makes it relatively more favorable to save in mortgages as compared to bonds.

I proceed by performing a counterfactual experiment in which the available mortgage contract is a 30-year mortgage with *adjustable* rates. My findings show that when mortgages have variable rates, the aggregate immediate response in consumption to the real interest rate shock is more than *six* times as large, as compared to when fixed-rate mortgages are used (0.58 percent vs 0.09 percent). The immediate increase in house prices is also substantially larger under adjustable-rate mortgages (2.59 percent as compared to 0.25 percent).

There are two distinct features of adjustable-rate mortgages that contribute to the stronger response. First, *all* mortgagors are affected when the mortgage interest rates change, as opposed to only those who buy a new house or refinance their mortgage under fixed-rate mortgages (FRMs). Second, the short-term mortgage interest rates of adjustable-rate mortgages (ARMs) respond more to temporary changes in the policy rate as compared to the long-term mortgage interest rates of FRMs. Given the 30-year repayment plans of mortgages, the temporary significantly lower mortgage interest rate under ARMs benefits mortgagors more than the smaller but more persistent decrease under FRMs. Furthermore, the sharper decline in the mortgage interest rate alleviates the payment-to-income constraint in the mortgage market.

Under ARMs, the change in the mortgage interest rate accounts for about half of the amplification in aggregate consumption relative to the economy with FRMs. The remaining 50 percent is attributed to the interaction effect of changes in both the mortgage interest rate and house prices. As previously discussed, the endogenous increase in house prices relaxes the LTV constraint among refinancers and enables larger cash-outs. Perhaps surprisingly, under ARMs, the share of refinancers increases more than under FRMs, even though current homeowners do not need to refinance to capture the lower interest rate under ARMs. Thus, the higher prevalence of refinancing is solely due to an increased demand for cashing-out housing equity.

The increase in the mortgage balances of refinancers contributes to an overall shift in savings away from mortgages into bonds, where the aggregate balances in both mortgages and bonds increase following the interest-rate reduction. This stands in sharp contrast to the economy with FRMs, where the savings are reallocated to mortgages instead of liquid bonds. The average LTV under ARMs increases following the reduction in the real interest rate, which is due to the higher share of refinancers who take up larger mortgages.

The key driver of the aggregate response in consumption, in particular under ARMs, is the improved consumption smoothing by constrained homeowners who refinance their mortgage. The importance of this mortgage channel of monetary policy has implications for a well-known puzzle in the monetary-policy literature: the forward guidance puzzle. A change in interest rates far into the future does little for currently constrained homeowners. Thus, the forward-guidance critique does not apply to this mortgage channel of monetary policy transmission.

Overall, my findings suggest that mortgages and housing play an important role in the transmission of monetary policy and the structure of the mortgage market impacts how effective monetary policy is. There is a significantly stronger response in consumption to a monetary policy shock when mortgages have adjustable rates, where the endogenous response in house prices is crucial for the amplification. The greater effectiveness follows from the behavior of mortgagors: consumption smoothing is improved among constrained households under ARMs. I conclude that including housing and mortgages in the analysis of monetary policy has qualitative implications for the transmission channels, and can have quantitatively important implications for aggregate responses. Thus, a detailed understanding of the mortgage market is a valuable input in monetary-policy analysis.

As already alluded to, there are several recent studies that investigate the role of mortgages in the transmission of monetary policy empirically.

Overall, this literature suggests that mortgagors are important for the transmission of monetary policy. In particular, households who experience changes in their mortgage interest payments adjust their consumption to a greater extent than other homeowners (Di Maggio et al., 2017, Flodén et al., 2019). Moreover, in countries where variable-rate mortgages are more common, house prices and consumption respond more strongly to monetary policy shocks (Calza et al., 2013). Cloyne et al. (2019) conclude that the aggregate response to monetary policy is largely driven by mortgagors and households with little liquid wealth. I view these empirical results as motivating facts that I rationalize in my model.

There is an extensive literature that studies the transmission of monetary policy within the framework of dynamic stochastic general equilibrium models. Recently, the importance of incorporating heterogeneous agents with various degrees of liquid and illiquid wealth has been emphasized by, for example, Auclert (2019) and Kaplan et al. (2018). In my model, housing is an illiquid asset and the costs associated with accessing housing equity are measured in the data. Greenwald (2018) and Hedlund et al. (2019) incorporate housing and mortgages in large structural models and find that endogenous changes in house prices amplify aggregate responses to monetary policy shocks, something that I also find.

Related to my finding that refinancing activities increase following an expansionary monetary policy shock, Chen et al. (2013) document that refinancing is negatively related to the business cycle. Moreover, in a simulation of the Great Recession, they find that depressed house values led to less refinancing. Beraja et al. (2018) show that the prevalence of mortgage refinancing is linked to house price growth, which in turn affects the spending responses to monetary policy. Eichenbaum et al. (2018) emphasize that the distribution of savings from refinancing is a key determinant of the efficacy of monetary policy. Hence, this strand of the literature shows that changes in house prices affect refinancing activities, and that refinancing is an important transmission mechanism of monetary policy. My results are in line with these findings, and I show that the extent to which monetary policy affects house prices varies substantially with the contract structure in the mortgage market.

The two papers that are the most closely related to mine are Garriga et al. (2017) and Wong (2019), which also investigate monetary policy under an FRM versus an ARM regime. Garriga et al. (2017) show that ARMs contribute to stronger responses of interest-rate shocks as compared to when FRMs are used, but do so in a structural model with two representative agents. My findings suggest that the distribution of constrained homeowners is important when quantifying the effects of changes in interest rates. Wong (2019) uses a life-cycle model to study the importance of refinancing behavior in the transmission of monetary policy. She also finds that variable-rate mortgages increase the aggregate response of consumption as compared to when FRMs are used. I endogenize the response in house prices in these two environments, where I find that under ARMs, house prices respond much more strongly to an interest rate shock, a finding that is empirically supported; see Calza et al. (2013). Further, I show that the stronger response in house prices under ARMs contributes substantially to the amplified response in aggregate consumption.

Overall, I contribute to the literature on the role of mortgages and housing in the transmission of monetary policy by quantifying the importance of changes in mortgage interest rates and house prices. I confirm empirical findings that monetary policy is more efficient in stimulating aggregate consumption under adjustable-rate mortgages as compared to when mortgages have fixed rates. Further, I highlight the mechanism through which the interaction between changes in house prices and mortgage interest rates leads to an amplification in aggregate consumption following an expansionary interest rate shock.

The paper is organized as follows. Section 2.2 describes the model. Section 2.3 proceeds by calibrating the model to U.S. data, and I compare the model to the data along a range of relevant variables. In Section 2.4, I show the impulse response functions from an unexpected real interest rate shock, and decompose the role of changes in mortgage interest rates and house prices. I proceed by analyzing how the transmission of monetary policy is affected if ARMs are used instead of FRMs. Section 2.5 concludes the paper.

2.2 Model

To study the consumption implications of changes in house prices and mortgage interest rates from a real interest rate shock, I use a heterogeneous agent life-cycle model with a detailed modeling of mortgage contracts and the housing market.[4] The setting represents a small open economy in which the interest rate is exogenous but where house prices and rental rates are equilibrium objects. In a given period, households choose to rent or own a house, the home size, the use of mortgage financing, savings, and consumption. House purchases are subject to transaction costs, and mortgage financing is restricted by down-payment and payment-to-income requirements. Furthermore, refinancing costs reduce the liquidity of housing equity. Mortgages are modeled as the most commonly used type in the U.S., i.e., long-term fixed-payment contracts. In this modeling environment, I am able to analyze the direct effects on demand of shocks to the real interest rate through changes in debt service costs, returns on liquid savings, and house prices.[5]

2.2.1 Households

The model is in discrete time. Households enter the economy at age $j = 1$, which represents the first period of working life, and work until age J^{ret}. When each household i is born, it receives an initial endowment of assets a, as in Kaplan and Violante (2014), and is allocated a permanent lifetime earnings state. In each period before retirement, the household is endowed with earnings y that depend on the individual lifetime earnings state and that are subject to idiosyncratic permanent and transitory shocks. Following retirement, households receive retirement benefits in a fixed proportion R of the permanent earnings in the last period of working life, subject to a cap. Households face an age-dependent probability of surviving to the next period $\phi_j \in [0, 1]$, and can live for a maximum of J periods.

There are three assets in the economy: owned housing h, mortgages m, and risk-free bonds b. Households realize utility from consumption c and housing services s, through a CRRA utility function with a Cobb-Douglas

[4] The model shares many features with Karlman et al. (2020).

[5] General equilibrium effects of demand on wages and profits, and changes in inflation, are not taken into account. Similarly, I investigate direct effects of changes in house prices, but I disregard general equilibrium effects of those on wages, for example.

aggregator over consumption and housing services

$$U_j(c, s) = e_j \frac{\left(c^{\alpha} s^{1-\alpha}\right)^{1-\sigma}}{1-\sigma}. \tag{2.1}$$

The age-dependent parameter e_j is a utility shifter that accounts for changes in household size over the life cycle (see, e.g., Kaplan et al., 2017). Housing services can be rented at a unit price p_r or attained by owning a house that is purchased at a unit price p_h. If a household chooses to own a home of size h, there is a linear transformation of owned housing into housing services such that $s = h$.

Households derive warm-glow utility from bequests, similar to in De Nardi (2004).

$$U^B(q) = \upsilon \frac{\left(q + \bar{q}\right)^{1-\sigma}}{1-\sigma}, \tag{2.2}$$

where υ denotes the weight that is attached to the utility from bequests, and $\bar{q}$ is a positive parameter that determines to what degree bequests are a luxury good. The amount of bequests q is given by the net worth of a household.

The illiquid nature of owned housing is characterized by transaction costs for both buying and selling a house, ς^b and ς^s, respectively. These are modeled as constant shares of the house value. Further, a homeowner needs to pay a periodic maintenance cost δ^h, also proportional to the house value. Mortgages are available to all homeowners in terms of fixed-payment long-term contracts. It is possible to refinance a mortgage, but it is subject to refinancing costs. The length of the available mortgage contract is indicated by l, and the number of periods left on a mortgage is given by $N = \min(J - j, l - ma)$, where ma is the mortgage age. I thus assume that mortgages have to be repaid in full by the age of certain death.[6] The minimum required mortgage payment is an age and mortgage-age

[6]This modeling choice is motivated by the fact that retirees tend to hold little debt and the terms of long-term mortgage contracts that are offered to retirees are often less favorable than those offered to working-age households.

dependent fraction $\chi_{j,ma}$ of the current mortgage balance

$$\chi_{j,ma}\, m = \frac{r_m(1+r_m)^N}{(1+r_m)^N - 1} m, \qquad \text{for } r_m > 0. \tag{2.3}$$

In steady state, the mortgage interest rate r_m is given by the risk free rate r plus an exogenous credit spread κ, i.e., $r_m = r + \kappa$. How the mortgage interest rate is affected by monetary policy depends on if the mortgage contracts have fixed or adjustable rates, which will be discussed in more detail in Section 2.4. New mortgage financing is restricted by a loan-to-value (LTV) requirement as well as a payment-to-income (PTI) cap. The loan-to-value constraint is given by

$$m' \leq (1-\theta) p_h h', \tag{2.4}$$

where θ specifies the required down-payment share of the house value $p_h h'$, and where prime indicates the current period choice of a state variable. The PTI requirement is modeled as

$$\frac{\chi_{j+1,ma} m' + (\tau^h + \varsigma^I) p_h h'}{n} \leq \psi, \tag{2.5}$$

where τ^h and ς^I represent property tax and home insurance payments, respectively, and n is permanent income.[7] Thus, ψ sets the maximum share of current permanent income that can be allocated to housing-related payments. These constraints need to be obeyed whenever a house is purchased or if a household chooses to refinance. In the latter case, the household has to pay a fixed refinancing cost ς^r, and a refinancing cost ς^r_p proportional to the mortgage size. A homeowner who does not refinance its mortgage needs to adhere to the minimum payment schedule

$$m' \leq (1+r_m)m - \chi_{j,ma}\, m. \tag{2.6}$$

In a given period, the state variable cash-on-hand x of a household is

[7]When banks evaluate the payment capabilities of prospective mortgage holders, three main components include mortgage payments, property taxes, and home insurance costs. Home insurance costs are only included for calibration purposes of the PTI requirement, see Section 2.3.1, and are not included in the households' budget constraint.

defined as follows,

$$x \equiv \begin{cases} y + (1+r)b - (1+r^m)m + (1-\varsigma^s)p_h h - \delta^h h - \Gamma & \text{if } j > 1 \\ y - \Gamma + a & \text{if } j = 1. \end{cases} \tag{2.7}$$

It consists of labor income or social security benefits y, any savings from liquid bonds less the mortgage balance including interest, the value of the house net of transaction costs, less property taxes and total tax payments Γ.[8] A household of the newborn cohort enters the economy with initial assets a. The total tax payments are made up by five different taxes

$$\Gamma \equiv \tau^l y + \mathbb{I}^w \tau^{ss} y + \tau^c rb + \tau^h p_h h + T(\tilde{y}). \tag{2.8}$$

A household pays local taxes on earnings given by the proportional tax rate τ^l.[9] All working-age households, as indicated by $\mathbb{I}^w$, also pay a social security tax τ^{ss}, proportional to earnings. Further, there is a capital income tax τ^c that applies to all earned interest, and the property tax τ^h is paid by homeowners as a share of their house value. Finally, $T(\tilde{y})$ captures the progressive federal labor income tax, where T is a non-linear function that takes taxable labor income after deductions $\tilde{y}$ as its argument. A household may deduct its mortgage interest payments, property taxes, and local labor income taxes. The federal income tax system is described in more detail in Section 2.2.3.

Let R, B, Ref, S denote the mutually exclusive and exhaustive cases where a household rents, buys a house, is a homeowner that refinances its mortgage, or is a homeowner that stays in its house and fulfills the minimum mortgage payment requirement, respectively. The dynamic household problem is described by the following Bellman equation where households discount future periods exponentially, with a discount factor β. Let $\mathbf{z} \equiv (h, m, ma, n, x)$, then for each $k \in \{R, B, Ref, S\}$,

$$V_j^k(\mathbf{z}) = \max_{c,s,h',m',b'} U_j(c,s) + (1-\phi_j)U^B(q') + \beta\phi_j \mathbb{E}_j \left[V_{j+1}(\mathbf{z}') \right]$$

[8] The definition of cash-on-hand includes the net revenue from selling a house. This is only included for computational simplicity, and a household that stays in its house does not incur a transaction cost.

[9] Local labor income taxes are deductible, and are included in the model to ensure that high-earning households benefit more from using itemized deductions.

subject to

$$c + b' + \mathbb{I}^R p_r s + \mathbb{I}^B (1+\varsigma^b) p_h h' + \mathbb{I}^{Ref,S}(1-\varsigma^s) p_h h + \mathbb{I}^{Ref}(\varsigma^r + \varsigma_p^r m') \leq x + m' \quad (2.9)$$

$$q' = b' + p_h h' - m' \quad (2.10)$$

$$s = h' \quad \text{if } h' > 0 \quad (2.11)$$

$$m' \geq 0 \quad \text{if } h' > 0 \quad (2.12)$$

$$m' = 0 \quad \text{if } h' = 0 \quad (2.13)$$

$$c > 0, s \in S, h' \in H, b' \geq 0,$$

where $\mathbb{I}^k$ are indicator variables that take the value of one for the relevant case and zero otherwise.[10] Equation (2.9) specifies the household's budget constraint, and equation (2.10) defines the bequests. The last four rows state a set of constraints including that a homeowner may not be a landlord and mortgages may only be used to finance owned housing. A household that buys a house or refinances its mortgage also needs to fulfill the LTV and PTI requirements specified in equations (2.4) and (2.5), and a homeowner that stays in the same house but does not refinance its mortgage needs to fulfill the minimum mortgage payment requirement in equation (2.6). Additionally, rented housing services are only available in discrete sizes contained in the ordered set $S = \{\underline{s}, s_2, s_3, ..., \overline{s}\}$. Owned housing is limited to a set H, which is a proper subset of S. Specifically, the smallest size $\underline{h}$ in H is larger than the smallest size in S, and above and including $\underline{h}$ the two sets are identical.[11] The solution to the household problem is given by

$$V_j(\mathbf{z}) = \max_k V_j^k(\mathbf{z}) \quad (2.14)$$

[10]To ensure that bequests cannot be negative, the utility from bequests is not discounted, but the parameters of the bequest function are estimated to match moments in the data.

[11]It is common in the literature to restrict the minimum house size available for owning, e.g., see Cho and Francis (2011), Floetotto et al. (2016), Gervais (2002), and Sommer and Sullivan (2018).

for $k \in \{R, B, Ref, S\}$, with the corresponding set of policy functions

$$\left\{c_j(\mathbf{z}), s_j(\mathbf{z}), h'_j(\mathbf{z}), m'_j(\mathbf{z}), b'_j(\mathbf{z})\right\}.$$

2.2.2 Rental market

The rental market consists of a unit mass of homogeneous rental firms f that provide rental housing to households. Firms operate in a competitive market and are owned by foreign investors. The required rate of return of the investors equals the after-tax return on bonds. In steady state, the house price is constant, i.e., $p_h = p'_h$, and the equilibrium rental rate p_r^{ss} per unit of housing is given by the following user-cost-formula,

$$p_r^{ss} = \left[1 - \beta_f + \beta_f \left(\delta^r + \tau^h\right)\right] p_h, \tag{2.15}$$

where $\beta_f = \frac{1}{1+(1-\tau^c)r}$ is the investors' discount factor. Thus, the rental rate is such that, after paying maintenance costs and property taxes, rental firms earn their required rate of return. Both the maintenance cost and the property taxes are given by constant shares of the rental property value in the next period. The maintenance cost covers the depreciation of rental property $\delta^r p_h$, where $\delta^r > \delta^h$.[12]

Motivated by the finding that rental rates often adjust slowly to changes in house prices, I assume that owners of rental firms have a long-term investment horizon. Rental firms own the steady-state stock of rental housing and in any period where the demand for rental housing deviates from the steady-state demand, rental firms transact in the housing market such that their rental stock equals demand. The present value of the accounting profits in steady state consists of the rental revenue less the discounted costs in the next period,

$$\pi_f^{ss} = p_r^{ss}\bar{S} - \beta_f \left(\delta^r + \tau^h\right) p_h \bar{S},$$

where $\bar{S}$ is the steady-state rental stock. Let S denote the demand for rental housing in any given period. The present value of the accounting

[12]The assumption that the depreciation rate is higher for rental property than for owned housing is common in the literature (see, e.g., Piazzesi and Schneider, 2016), and is supported by the potential moral hazard problem in rental housing markets.

profits in a period with rental demand S is then given by

$$\pi_f^{tr} = p_r S - \beta_f(\delta^r + \tau^h)p_h' S + p_h(\bar{S} - S) - \beta_f p_h'(\bar{S} - S),$$

where p_r is the rental rate, p_h is the house price in the current period, while p_h' is the house price in the next period, and the stock $\bar{S} - S$ of housing is transacted in the market. For rental firms to earn the same return on their investments as in steady state, i.e., $\pi_f^{tr} = \pi_f^{ss}$, the rental rate p_r is provided by

$$p_r = (1 - \beta_f)p_h + \beta_f(\delta_r + \tau_h)p_h' + \beta_f \Delta p_h' \frac{S - \bar{S}}{S}, \tag{2.16}$$

where $\Delta p_h' \equiv p_h - p_h'$. Thus, the rental rate is such that the investors earn their required rate of return, after paying maintenance cost and property taxes, and after accounting for potential fluctuations in the house price on the share of the rental stock that is transacted in the market.

2.2.3 Government

The government in the model has two main tasks: providing retirement benefits to households and taxing the agents in a manner that reflects the U.S. tax code. Overall, the government runs a surplus, which it spends wastefully, or on matters that do not affect the agents in the model. Rental firms pay two taxes, property taxes in proportion to the value of the rental stock and capital income taxes on their profits. As discussed in Section 2.2.1, households pay five different taxes. Working-age households pay social security taxes, and all households pay local and federal labor income taxes. Additionally, there is a tax on earned interest on savings, and homeowners pay a property tax.

To capture the level of progressivity in the U.S. federal income tax schedule, I use a continuous and convex tax function as in Heathcote et al. (2017), where the argument is taxable earnings net of deductions $\tilde{y}$. The function is given by

$$T(\tilde{y}) = \tilde{y} - \lambda \tilde{y}^{1-\tau^p}, \tag{2.17}$$

where parameters λ and τ^p control the level and the degree of progressivity in the tax system.

Taxable earnings are determined by labor income or retirement benefits less any deductions. Working-age households can choose to use an itemized deduction, a standard deduction, or not deduct anything, while retired households may only choose between the latter two. If a household chooses to use itemized deductions, it can deduct mortgage interest payments, property taxes, and local labor income taxes. The most favorable type of deduction depends on a household's earnings and the size of any payments that are deductible under the itemized specification. Specifically, a household chooses the type of deduction that minimizes $T(\tilde{y})$, subject to

$$\tilde{y} \in \begin{cases} \{\max(y - ID, 0), \max(y - SD, 0), y\} & \text{if } j \leq J_{ret} \text{ and } ID > SD \\ \{\max(y - SD, 0), y\} & \text{otherwise} \end{cases} \tag{2.18}$$
$$\text{where } ID = r^m m + \tau^h p_h h + \tau^l y.$$

ID denotes the deductible amount if a household uses itemized deductions, and SD is the tax subsidy available to households that opt for the standard deduction.

To summarize, I include the main components of the U.S. tax system related to housing and mortgages in the model, i.e., imputed rents are not taxed, property taxes and mortgage interest payments are deductible, both itemized and standard deductions are available to households, and the earnings tax is progressive.

2.2.4 Solving the model

The dynamic programming problem is solved recursively. The steady state of the baseline economy is solved for by computing the value and policy functions, and simulating an economy where households behave according to the solved for decision rules. The state space and the transitory earnings shocks are discretized to solve the model. The equilibrium house price is set exogenously, and the rental rate is then given by equation (2.15). The steady-state total demand for housing, both rental and owned housing, provides the total supply of housing, which is held constant throughout the analysis.

To analyze the effects of an interest rate shock, I solve for a transitional equilibrium from an unexpected shock to the real interest rate. Given the path of the real interest rate, I compute the transition path of the mortgage interest rate. For fixed-payment mortgage contracts with long-term interest rates, I assume that the mortgage interest rate is given by the geometric mean of the expected gross periodic mortgage interest rates, for the lifetime of the mortgage. When there is a change in the mortgage interest rate, the repayment plans adjust accordingly for all households who take up a new mortgage, i.e., for those who take up a mortgage when buying a new house or when refinancing. For the analysis where mortgages have variable rates, the mortgage interest rate at any point in time is given by the periodic risk-free interest rate plus the credit spread κ. With adjustable-rate mortgages, the repayment plans update for all new and outstanding mortgages, to capture the change in the mortgage interest rate.

For the transitional equilibrium, a vector of house prices and a vector of total rental housing supply are solved for, such that in each period of the transition, the total demand for housing, both rental and owned housing, equals the total supply, and the demand for rental housing equals the rental supply, given the rental rate in equation (2.16). I assume that households have perfect foresight of the transition paths of the interest rates and the house and rental prices. The equilibrium definitions are stated in Appendix 2.A, and a more detailed description of the solution method is provided in Appendix 2.B.

2.3 Calibration

The model is parameterized to the U.S. economy in 1989 to 2013. I choose to use average data moments across many years in an attempt to avoid cyclicalities and capture a steady state of the economy. Housing wealth and household debt have varied substantially over time, for example, and the goal of the model is to investigate real interest rate shocks that hit an economy that is in a steady state. Most of the parameter values are chosen from data or other studies. The remaining parameters are estimated by jointly minimizing the distance between several relevant equilibrium moments in the model and their data counterparts. A model period corresponds to one year.[13]

2.3.1 External model parameters

The independently calibrated parameters are listed in Table 2.1.

Parameter	Description	Value
σ	Coefficient of relative risk aversion	2
r	Interest rate	0.03
κ	Yearly spread, mortgages	0.014
τ^l	Local labor income tax	0.05
τ^c	Capital income tax	0.15
τ^{ss}	Payroll tax	0.153
τ^h	Property tax	0.01
l	Mortgage contract length	30
θ	Down-payment requirement	0.20
ψ	Payment-to-income requirement	0.28
δ^h	Depreciation, owner-occupied housing	0.03
ς^I	Home insurance	0.005
ς^b	Transaction cost if buying house	0.025
ς^s	Transaction cost if selling house	0.07
ς^r_p	Proportional refinancing cost	0.01
R	Replacement rate for retirees	0.50
B^{max}	Maximum benefit during retirement	0.61

Table 2.1: Independently calibrated parameters, taken from the data or other studies
Note: The table lists calibrated parameter values, and where relevant, these are annual.

Demographics

Households enter the economy at age 23 and work until age 65. The probability of dying at any age $(1 - \phi_j)$ is set to match the observed and projected mortality rates for males born in 1950, in the Life Tables for the U.S., social security area 1900-2100 (see Bell and Miller (2005)). The maximum age J in the model is 83 years. The age-dependent equivalence scale parameters e_j are determined from the Panel Study of Income Dynamics (PSID). The parameter values are set to the square root of

[13]In an earlier version of this paper the model period length corresponded to three years. The calibration and the main results from the model with a period length of three years are found in Appendix 2.E. The main findings are robust to changing the frequency of the model.

the predicted values from a regression of family size on a third-order polynomial of age.

Preferences and interest rates

The parameter governing households' relative risk aversion σ is set to 2, which gives an elasticity of intertemporal substitution of 0.5. The real interest rate on risk-free bonds r is set to 0.03. This is consistent with the average yield on 30-year constant maturity nominal Treasury securities, deflated by the yearly headline Consumer Price Index (CPI). Between 1997 to 2013, this average real rate was 0.034 (Federal Reserve Statistics Release, H15, and the Bureau of Labor Statistics, Databases & Tables, Inflation & Prices). The mortgage spread κ is set to 0.014. This is given by the average yearly difference between the rate on 30-year fixed-rate conventional home mortgage commitments and the above nominal Treasuries, from 1997 to 2013. Thus, the steady-state mortgage interest rate is 0.044.

Taxes

The local labor income tax rate is determined by the average state and local labor income tax rate for households that itemize deductions, which was 5 percent in 2011 (Lowry, 2014). The tax rate on capital income is chosen to be the maximum rate that applies to long-term capital income for most taxpayers, which is 15 percent. The social security tax paid by the working age population, i.e., the payroll tax, is set to 15.3 percent of earnings. This rate captures the payroll taxes that are paid by both employees and employers (Harris, 2005). The property tax varies significantly across U.S. states. I choose a property tax rate of 1 percent, which is approximately the median rate in the American Housing Survey (AHS) for the 2009, 2011, and 2013 survey years.

Housing and mortgage markets

In the baseline economy, mortgages are modeled as 30-year fixed-payment mortgages, which is the most commonly used mortgage contract in the U.S. Hence, l is set to 30. In the exercise where I compare the baseline economy to a setting where mortgages have adjustable rates, mortgages with variable rates still need to be paid off in 30 years. The minimum down-payment requirement θ in the model is 0.20. Below this threshold, mortgage lenders often require an extra insurance. In the period leading up to the Great Recession, it became more common to borrow above 80 percent of the house value, but this period can be seen as an outlier. Between 1978 and 1992, the average down payment of first-time buyers in the U.S. ranged from 11.4 to 20.5 percent of the house value (U.S. Bureau of the Census, Statistical Abstract of the United States (GPO), 1987, 1988, and 1994). The payment-to-income requirement ψ is set to 0.28, as in Greenwald (2018).

The depreciation rate on owner-occupied housing δ^h is taken from Harding et al. (2007) and is set to 0.03. This value is the estimated median depreciation rate, gross of maintenance. The home insurance rate ς^I is equal to 0.005 of the house value. This figure is taken from the AHS, where the median property insurance payments correspond to approximately half of the median property tax payments.

The transaction costs for buying and selling a house, ς^b and ς^s, are set to 2.5 and 7 percent, respectively. These numbers are taken from Gruber and Martin (2003) who use median transaction costs in CES data to estimate the transaction costs in proportion to the house value. The refinancing cost that is proportional to the mortgage size ς^r_p is set to 0.01, as in Gorea and Midrigan (2017).

Assets of newborns

The newborn households in the model are endowed with the initial assets a conditional on earnings. The allocation is based on the method in Kaplan and Violante (2014). In the Survey of Consumer Finances (SCF), households of age 23-25 are divided into 21 groups based on earnings. Within each group, the share of households with asset holdings above 1,000 2013 dollars is calculated, along with their median asset values. The median asset holdings are then scaled by the median earnings of households aged 23-64. Within each of the comparable 21 groups in the model, ranked on initial earnings, the shares found in the SCF divide the households into low-earners who do not receive any initial assets, and high-earners who are allocated the median asset value consistent with that group, and rescaled by the median earnings of working-age households in the model.

Labor income and social security

In each period, households are endowed with exogenous earnings. The estimation of the earnings process follows Cocco et al. (2005). There is a deterministic life-cycle component of labor income, and in each period during working age, households' earnings are subject to idiosyncratic permanent and transitory shocks. For household i, of age $j \leq J^{ret}$, the log of labor income is given by

$$\log(y_{ij}) = \alpha_i + g(j) + \eta_{ij} + \nu_{ij} \quad \text{for } j \leq J_{ret}, \tag{2.19}$$

where α_i is a household fixed effect with the distribution $N(-\frac{\sigma_\alpha^2}{2}, \sigma_\alpha^2)$. The function $g(j)$ captures the deterministic life-cycle component of earnings, while η_{ij} and ν_{ij} are the permanent and transitory components, respectively. The transitory earnings shock ν_{ij} is i.i.d., with the distribution $N(-\frac{\sigma_\nu^2}{2}, \sigma_\nu^2)$. The permanent earnings risk is modeled as a random walk, where there are i.i.d. shocks ζ_{ij} with the distribution $N(-\frac{\sigma_\zeta^2}{2}, \sigma_\zeta^2)$, such that

$$\eta_{ij} = \eta_{i,j-1} + \zeta_{ij} \quad \text{for } j \leq J_{ret}. \tag{2.20}$$

In the model, the permanent earnings state n_{ij} consists of the three permanent components of labor income, i.e., $\log(n_{ij}) = \alpha_i + g(j) + \eta_{ij}$. In retirement, households receive a constant fraction R of permanent earnings in the last period of working life, subject to a cap B^{max}. Thus, there is no labor-income uncertainty in retirement.

$$\log(y_{ij}) = \min\left(\log(R) + \log(n_{i,J_{ret}}),\, \log(B^{max})\right) \quad \text{for } j \in]J_{ret}, J] \tag{2.21}$$

The labor income process is estimated using PSID data from 1970 to 1992. See Karlman et al. (2020) for a more detailed description of the data. A linear fixed-effect regression of the log of households' earnings on dummies for age, marital status, family composition, and education, is run to estimate the deterministic life-cycle profile. The components $g(j)$ are given by fitting a third-order polynomial to the mean predicted earnings by age from the regression. To estimate the variances of the permanent and transitory earnings shocks, I use a similar method as in

Carroll and Samwick (1997). The variance of the fixed-effect shock is found by computing the residual variance of earnings that is left after accounting for the life-cycle component and the estimated variances of the permanent and transitory shocks, for households of ages 23 to 25. The estimated variances are presented in Table 2.2.

Parameter	Description	Value
σ_α^2	Fixed effect	0.156
σ_ζ^2	Permanent	0.012
σ_ν^2	Transitory	0.061

Table 2.2: Estimated variances
Note: The three variances are the estimated variances for: the fixed-effect earnings shock that households realize when they enter the economy, and the permanent and transitory earnings shocks to which households are subject before retirement. Estimated using PSID data.

The replacement rate R for retirees is chosen to be 50 percent of earnings in the last period of working life, which is taken from Díaz and Luengo-Prado (2008). The maximum benefit limit B^{max} is computed from Social Security Administration (SSA) data, and is equal to 0.61 in the model. This number can be evaluated relative to the mean of expected annual earnings during working life that is normalized to one.

2.3.2 Estimated parameters

The parameters that I estimate through simulated method of moments are listed in Table 2.3. The parameters are estimated simultaneously, but the most relevant target moments for the respective parameters are listed in the table along with their values in the data and in the model.

Unless otherwise stated, the data moments are computed from the SCF, using pooled data over the 1989 to 2013 waves. The parameter α in the utility function controls the share of expenses that is allocated to consumption as opposed to housing services. The target moment that is used to discipline this parameter is the median house value-to-earnings, conditional on owning, which is an indicator of the relative importance of housing costs compared to other expenses. The discount factor β affects borrowing and savings decisions, and is therefore estimated by targeting the median LTV in the economy. The benefit of buying a house instead of renting is in the model affected by the preferential tax treatment of owned

Parameter	Description	Value	Target moment	Data	Model
α	Consumption weight	0.75	Median house value-to-earnings	2.30	2.30
β	Discount factor	0.92	Median LTV	0.35	0.35
δ^r	Depreciation rate, rentals	0.055	Homeownership rate, age < 35	0.44	0.40
$\underline{h}$	Min. owned house value	0.35	Homeownership rate	0.70	0.73
ς^r	Fixed refinancing cost	0.12	Refinance rate	0.08	0.08
$\bar{q}$	Luxury of bequests	6.8	Net worth p75/p25, age 68-76	5.37	5.26
υ	Utility shifter of bequests	190	Median net worth, age 75/50	1.44	0.68
SD	Standard deduction	0.081	Itemization rate	0.53	0.53
λ	Level, tax function	0.975	Average marginal tax rates	0.13	0.13
τ^p	Progressivity, tax function	0.17	Distr. of marginal tax rates	See text	

Table 2.3: Estimated parameters
Note: Estimated parameters using simulated method of moments. The resulting parameter values are shown in column three. Column five displays the relevant target moment value in the data, while column six shows the comparable moment value in the model when the listed parameter values are used. The values are annual when relevant. The minimum owned house size $\underline{h}$, the fixed refinancing cost, the luxury parameter in the utility function for bequests, and the standard deduction SD, can be evaluated relative to the mean of expected annual earnings during working life that is normalized to one.

housing as well as the difference between the depreciation rate of owned and rental housing. To estimate the depreciation rate of rental housing δ^r, I use the homeownership rate among households aged below 35 as a target moment. The overall homeownership rate is used to estimate the value of the smallest housing unit available to own $\underline{h}$. To account for the frictions in the mortgage market, I estimate the fixed refinancing cost ς^r. In the model in steady state, the interest rate is constant and thus there is no reason to refinance to capture changes in the mortgage interest rate. The fixed refinancing cost is therefore estimated by targeting the share of households that refinance while also extracting equity from the house. This data moment value is taken from Gorea and Midrigan (2017).

The two parameters of the utility function of bequests are disciplined by two target moments related to savings. First, the parameter that captures the extent to which bequests are a luxury good $\bar{q}$ is estimated by targeting the fraction of net worth in the 75th over the 25th percentile, for households aged 68 to 76. Second, the parameter that determines the weight that is assigned to the utility from bequest υ is calibrated to match the fraction of median net worth of households aged 75 over the median net worth of 50-year-olds. Finally, I estimate three parameters related to the tax system. The level of the standard deduction SD impacts to what extent households use the itemized deduction, which in turn influences how households are differently affected by a change in mortgage interest rates. The standard deduction amount is used to match the itemization rate among the working-age population. The parameter λ that influences the level of the tax and transfer function $T(\tilde{y})$ is estimated to match the average marginal tax rate in the economy; while the progressivity parameter τ^p is estimated to approximate the distribution of households across statutory federal labor income tax brackets. The latter is done by computing the shares of households that are exposed to the different tax brackets. In the model, where the federal labor income tax rate is continuous, households are allocated to their nearest statutory bracket. I solve for the τ^p that minimizes the sum of the absolute values of the difference in shares in the model versus in the data. The data on the shares and the average marginal tax rate are taken from the Congressional Budget Office in 2005 (Harris, 2005), and the tax rates for the brackets correspond to the tax code from 2003 to 2012 (The Tax Foundation, 2013).

2.3.3 Model versus data

To evaluate how well the model reflects the data along dimensions that are not targeted in the estimation, I present a comparison between the model and the data for moments that are particularly important for how households respond to interest-rate changes. The strength of the mortgage cash-flow channel of monetary policy depends on the types of households that are homeowners and mortgagors, and how large mortgages different households use. In Figure 2.1, the life-cycle profiles of homeownership, median LTV, and median mortgage and housing to earnings are presented. The life-cycle patterns are clear: young homeowners are the most in debt and have the largest mortgage balances relative to earnings. The model successfully matches the life-cycle profiles computed from the SCF, with the exception of homeownership, where too many middle-aged households and too few old households are homeowners.

The prevalence of constrained households further impacts the importance of the mortgage cash-flow channel. A comparison of the distributions of liquid asset-to-earnings, LTV, and net worth-to-earnings, in the model versus the data is displayed in Figure 2.2, along with a correlation plot of leverage and liquid assets.[14]

[14] I define liquid assets as checking, savings, money market, and call accounts, prepaid

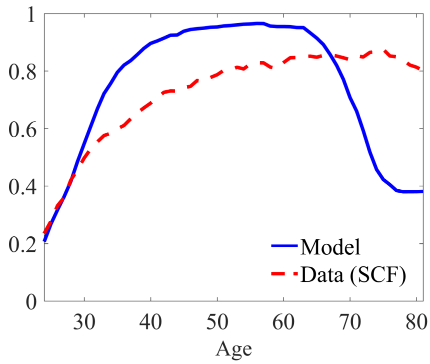

(a) Homeownership rate

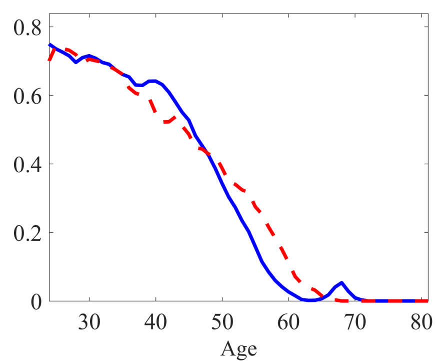

(b) Median LTV

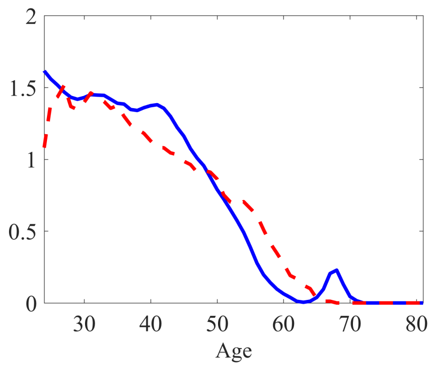

(c) Median mortgage-to-earnings

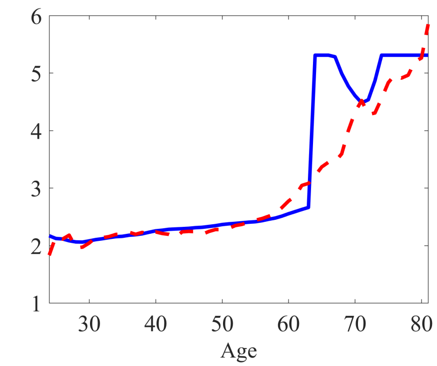

(d) Median house-to-earnings

Figure 2.1: Comparison between model and data: non-targeted life-cycle profiles
Note: Data refers to the Survey of Consumer Finances, for the survey years 1989 to 2013.

A household with a liquid-asset-to-earnings ratio of less than 0.5 is often referred to as a hand-to-mouth household in the literature. The model replicates fairly well the distribution of households with low liquid savings. The share of households with a liquid-asset-to-earnings ratio of less than 0.5 is 24 percent in the model, and 38 percent in the SCF. Among homeowners, i.e., the relatively wealthy households, this share is approximately 32 percent in the model and 31 percent in the data. Hence, the model does well in terms of matching the prevalence of wealthy

cards, cash, bonds and bills, less any credit card debt balance.

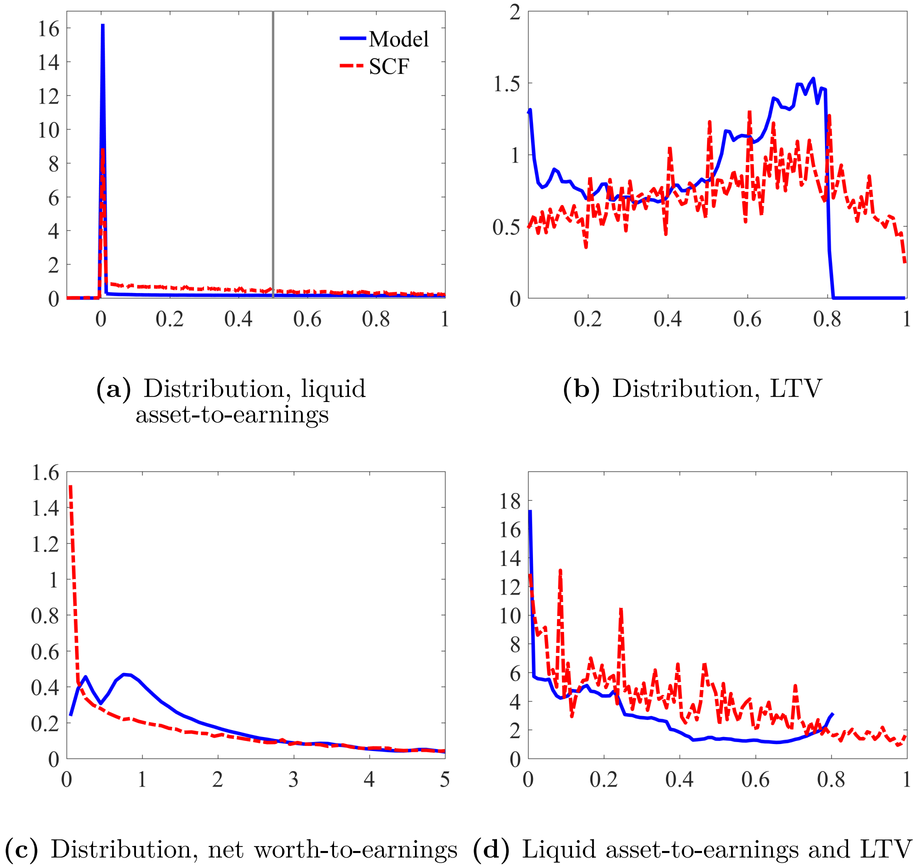

(a) Distribution, liquid asset-to-earnings

(b) Distribution, LTV

(c) Distribution, net worth-to-earnings

(d) Liquid asset-to-earnings and LTV

Figure 2.2: Comparison between model and data: distributions
Note: Data refers to the Survey of Consumer Finances, for the survey years 1989 to 2013. Values of liquid asset-to-earnings above 300 in the data are censored.

households with low liquid savings, but clearly underestimates the share of liquidity-constrained renters. The distributions of LTVs and net worth-to-earnings also match the data well. However, in the data there are households with LTVs above 0.8, which is the cap in the model. The extent to which homeowners with mortgages are constrained in their spending can be evaluated by the correlation between liquid asset-to-earnings and LTVs. The model shows the same pattern as that found in the data: more leveraged households tend to have less liquid savings.

2.4 Results

In this section, I first present impulse response functions (IRFs) of an unexpected real interest rate shock, and proceed by exploring the role of changes in mortgage interest rates and house prices for aggregate responses. In Section 2.4.2, I investigate the importance of the mortgage contract specification, i.e., the fixed-payment long-term mortgage. Specifically, I compare the IRFs in the baseline setting with fixed-rate mortgages (FRMs) to those when adjustable-rate mortgages (ARMs) are used.

In order to study the effects of a real interest rate shock, I use an empirically estimated path of the real interest rate from a shock of -100 basis points (bp) to the nominal interest rate. The estimated path of the real interest rate is the impulse response function from the identified Romer and Romer (2004) monetary policy shock in Auclert et al. (2020). The negative shock of 100bp to the nominal interest rate translates into an immediate reduction of the real interest rate of approximately 80bp.

I start from a steady state with an invariant distribution of households, and compute the non-linear IRFs to the "MIT shock" of the real interest rate. Following Boppart et al. (2018), these IRFs can be used to provide a linearized solution to the model with aggregate risk, i.e., only first-order effects of aggregate shocks are considered, as with standard first-order perturbations. The shock occurs just before the households make any decisions. There is an immediate adjustment of the paths of prices, but any cash-flow effects through changes in mortgage payments occur at the earliest one period after the shock is realized.

2.4.1 Fixed-rate mortgages and monetary policy

A shock to the real interest rate impacts mortgage interest rates. With FRMs, I assume that the mortgage interest rate of the long-term contract is given by the geometric mean of the expected gross yearly mortgage interest rates, for the lifetime of the mortgage.[15] In the given calibration, the contract length of a mortgage l is 30 years. Thus, the mortgage interest rate r_m at time t is the geometric mean of the expected gross yearly mortgage interest rates for the next 30 years. In Figure 2.3, the path of the mortgage interest rate along with the path of the real interest

[15] I assume that the credit spread κ remains constant over time.

rate on bonds are presented. With FRMs the currently available mortgage interest rate affects the mortgage interest payments of homeowners who take up a new mortgage, i.e., those who purchase a new home and use mortgage financing, and those who refinance an existing mortgage. A household that takes up a new mortgage in the period when the interest rate shock occurs receives a mortgage interest rate of 4.37 percent for the next 30 years, instead of the steady-state rate of 4.40 percent.

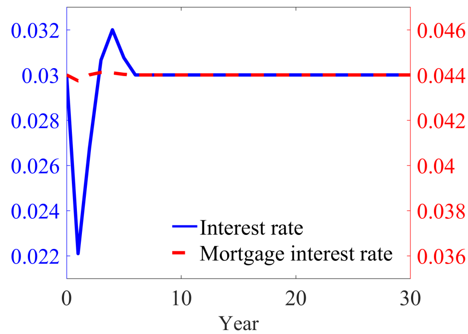

Figure 2.3: Long-term mortgage interest rates
Note: The paths of the real interest rate and the long-term mortgage interest rate. The paths follow an unexpected nominal interest rate shock of -100bp, where the path of the real interest rate corresponds to the estimated impulse response function in Auclert et al. (2020). The mortgage interest rate reads off the right-hand side y-axis.

The focus of this paper is to explore the direct demand effects of a shock to the real interest rate, and the role played by changes in the mortgage interest rates and house prices. I begin by investigating the IRF for consumption, which is presented in Figure 2.4. The immediate aggregate response of consumption to the expansionary real interest rate shock of 80bp is approximately a 0.09 percent increase. At a first glance, the IRF may look somewhat unorthodox. However, bear in mind that aggregate income is kept constant in the analysis, since the emphasis is placed on direct demand effects. As such, the main effects of a change in the real interest rate work through a reallocation of resources over time, and not through a change in the available resources. Thus, the general pattern of consumption, with the sharp drop following the initial increase, is mainly a result of intertemporal substitution. In a general-equilibrium analysis this drop in consumption would be counteracted by an endogenous response of households' earnings (through employment and wages). In

Appendix 2.D, I present a step towards such an analysis. In that analysis, I include an estimated earnings response, which is taken from the impulse response function for output in Auclert et al. (2020). I show that the mechanisms and the findings in the main analysis of this paper survive when there is a response in earnings. Further, I find that the decline in consumption is then more gradual.

As many studies have shown, an endogenous response in labor income following a monetary policy shock is an important part of transmission. Nonetheless, in order for there to be general equilibrium effects on earnings, there must first be direct effects on demand. The purpose of this paper is to specifically study how those short-run responses are affected by changes in cash-flows through the mortgage market, and inspect the mechanisms at work.

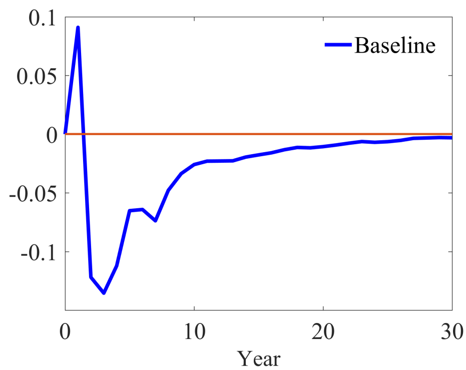

Figure 2.4: Impulse response function for consumption (%)
Note: The baseline model with fixed-rate mortgages. The impulse response function follows an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rate, as displayed in Figure 2.3.

There are three direct channels and one equilibrium channel through which the transmission of monetary policy works in the model with housing and mortgages. First, households are directly affected by changes in the real interest rate through the return on savings. When the savings rate is altered, households substitute consumption intertemporally by adjusting their level of savings, i.e., the traditional intertemporal-substitution channel. The second direct effect works through changes in cash-flows that, in turn, affect the level and timing of the lifetime resources of the agents. Specifically, the lower return on savings in risk-free bonds reduces the overall resources available over the lifetime. On the other hand, households who take up a new mortgage are affected through lower mortgage payments, which increase their lifetime resources. The third direct effect works through a relaxation of a credit constraint when the mortgage interest rate is reduced. In particular, it is easier for the marginal buyer or refinancer to fulfill the PTI requirement when the mortgage interest rate is lower, ceteris paribus. Finally, the fourth channel comprises the general equilibrium effect of changes in house prices and rental rates, which affects existing and potential homeowners, and renters, differently.

Changes in mortgage interest rates and house prices

To quantify the role of changes in mortgage interest rates and house prices, I compute the IRFs under the following assumptions i) Mortgage interest rates and house prices adjust endogenously; ii) House prices are constant; iii) Mortgage interest rates are constant; and iv) Both house prices and mortgage interest rates are constant. Figure 2.5 presents the IRFs for consumption under the different equilibrium assumptions. There is a slightly larger immediate response in consumption when the mortgage interest rate and house prices adjust as compared to when they are held constant.

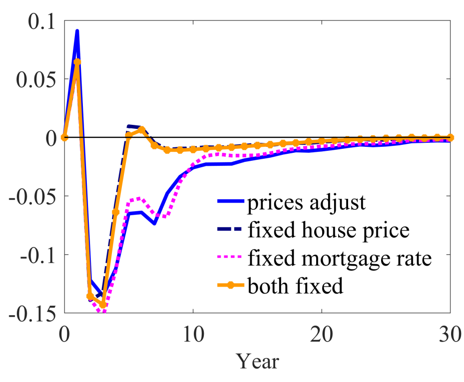

Figure 2.5: Impulse response functions for consumption (%)
Note: A decomposition of aggregate consumption responses under different equilibrium assumptions for house prices and the mortgage interest rate. The impulse response functions follow an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rate, as displayed in Figure 2.3.

To investigate the source of the amplification in aggregate consumption due to changes in the mortgage interest rate and house prices, I compute consumption responses for the four mutually exclusive categories of households in the model: house buyers, refinancers, stayers who follow the amortization plan, and renters. Table 2.4 presents the immediate mean consumption responses for these groups of households in the period when the interest rate shock occurs.[16] The groups are defined based on the tenure statuses of households in the period of the interest rate shock. The deviation in consumption is then computed as the difference in the mean consumption of a group of households in the period when the interest rate shock occurs as compared to the mean consumption of the same households in the steady state.

Table 2.4 reveals that households who refinance their mortgage are the main drivers of the amplification in aggregate consumption. However, the groups of buyers and renters also respond more strongly when prices adjust. It is also clear that the interaction of the lower mortgage interest rate and the endogenous response in house prices contributes to the stronger consumption response in the aggregate.

		Overall	Buyers	Refinancers	Stayers	Renters
FRM		0.09	0.03	0.37	0.02	0.31
	fixed p_h	0.06	-0.48	0.04	0.08	0.17
	fixed r_m	0.06	-0.18	0.09	0.05	0.18
	fixed p_h & r_m	0.06	-0.29	-0.07	0.08	0.17

Table 2.4: Consumption responses in the period when the interest rate shock occurs (%)
Note: A decomposition of mean consumption responses of buyers, refinancers, stayers, and renters, under different equilibrium assumptions for house prices and the mortgage interest rate. The deviations of consumption, in percent, are computed for the period when the real interest rate shock occurs. The separation into buyers, refinancers, stayers, and renters is based on the tenure choice in the period of the interest rate shock. The responses follow an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rate, as displayed in Figure 2.3.

Figure 2.6 presents how the share of refinancers changes in response to the interest rate shock, and well as the responses of liquid savings,

[16]At a first glance, the reduction in mean consumption among some groups of households may appear odd. However, when households change their discrete tenure choice in response to the expansionary interest rate shock, this may very well lead to a reduction in consumption among some groups. This is due to both transaction costs for households who change their tenure status as well as a change in the composition of households in the groups defined by tenure status.

mortgage balances, and the endogenous movement in house prices. Figure 2.6a shows that the refinancing activities rise only slightly, immediately after the interest rate shock, but increase more in the period after the shock occurs. Somewhat surprisingly, the larger prevalence of refinancing in the period after the shock occurs is mainly driven by intertemporal substitution, as this behavior is present also when mortgage interest rates and house prices are constant.

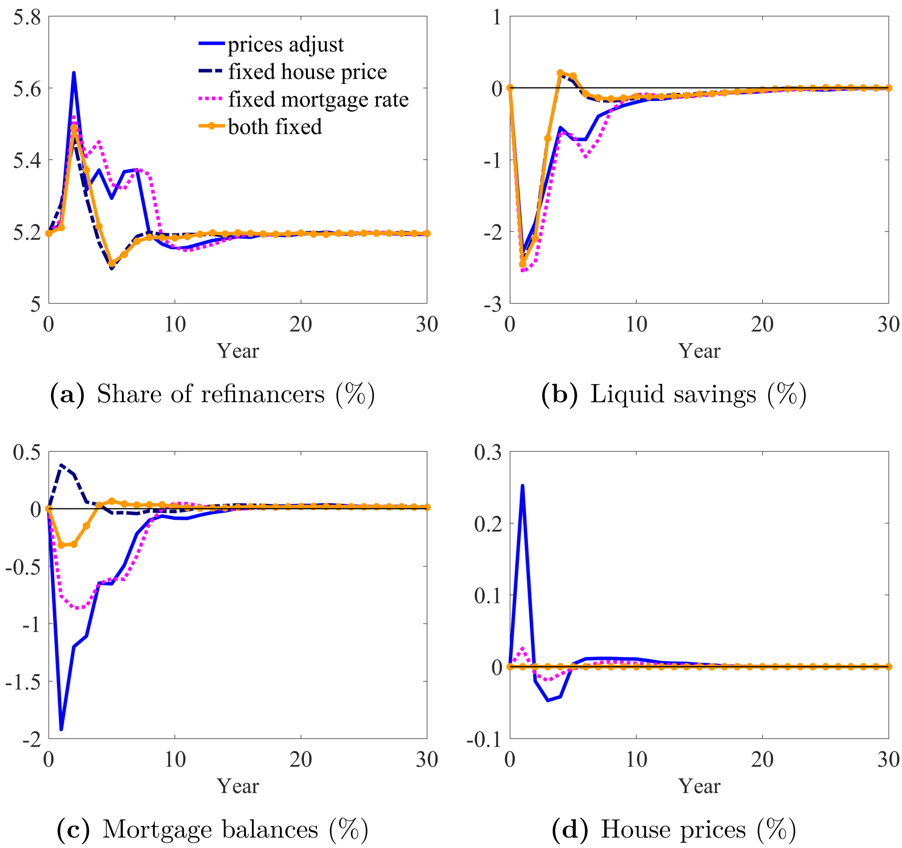

(a) Share of refinancers (%)

(b) Liquid savings (%)

(c) Mortgage balances (%)

(d) House prices (%)

Figure 2.6: Impulse response functions for refinancing behavior, savings in liquid bonds, mortgage balances, and house prices

Note: A decomposition of aggregate responses under different equilibrium assumptions for house prices and the mortgage interest rate. The impulse response functions follow an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rate, as displayed in Figure 2.3.

The immediately larger consumption response of refinancers stems both from the lower mortgage interest rate and from the endogenous increase in house prices, presented in Figure 2.6d.[17] The higher house prices enable larger cash-outs when refinancing a mortgage, given the LTV constraint. Households who find it optimal to refinance a mortgage do so in order to improve consumption smoothing. These homeowners have low levels of liquid savings, and most of their wealth is locked up in their house. Some of these households are liquidity constrained due to poor earnings realizations. Others are young homeowners who expect higher earnings in the future due to the upwards-sloping life-cycle profile of earnings, and therefore save in liquid bonds mainly for precautionary reasons. The long-term mortgage contracts tilt the mortgage payments relative to earnings such that the mortgage payments are more constraining early in life. The reduction in the mortgage interest rate and the higher house prices offer an opportunity to increase consumption today for these liquidity-constrained homeowners. The overall amplification in aggregate consumption is hence to a large extent driven by liquidity-constrained homeowners who refinance their mortgage and cash-out housing equity to smooth consumption.

All four channels through which monetary policy is transmitted contribute to the behavior of refinancers. The traditional intertemporal-substitution channel incentivizes a shift in consumption from future to earlier periods. The cash-flow channel, the PTI-channel, as well as the general-equilibrium channel of changes in house prices, allow for larger cash-outs among refinancing households.

First, let us consider the cash-flow channel, i.e., how households' behavior is affected by changes in their cash flows. As previously noted, there are no immediate cash-flow effects on savings and mortgage payments resulting from the interest rate shock. However, households are forward looking and take into account how their future cash flows are affected. Mortgagors who refinance receive a lower mortgage interest rate and benefit from lower future payments on their mortgage for an extended period of time. This reduction in payments can be seen as a persistent extra periodic income to these households. If households are unconstrained in their spending, such an extra income is smoothed over the life cycle. However, as discussed, households who refinance do so

[17] The corresponding figure for the rental rate is displayed in Appendix 2.C, Figure 2.13a.

exactly because they are in need of liquidity. With the knowledge of lower mortgage payments in the future, the cash-flow channel contributes to larger immediate consumption responses of the refinancers. The reduction in future mortgage interest payments enables a decrease in precautionary savings and a corresponding increase in consumption. This explains why refinancers increase consumption significantly also when only mortgage interest rates decrease and house prices are held fixed.

Second, the reduction in the mortgage interest rate relaxes the PTI constraint for households who refinance. One important component that determines what mortgage a household can afford is the expected mortgage interest payments, as seen in equations (2.5) and (2.3). The relaxed PTI constraint enables households to take up larger mortgages. As a result, households who refinance can improve consumption smoothing further.

Finally, the general equilibrium effect on house prices allows for larger cash-outs among those who refinance. Immediately after the interest rate shock occurs, house prices increase by approximately 0.25 percent. The higher house prices influence the liquidity of housing equity by affecting how far existing homeowners are from a binding LTV constraint. Households who use cash-out refinancing can extract a larger amount before being constrained by the LTV requirement. As seen in Figures 2.6a and 2.6d, the house-price increase in the period of the interest rate shock does not lead to a greater prevalence of refinancing but instead to larger cash-outs.

As previously noted, the groups of buyers and renters in the period of the interest rate shock also respond more strongly when house prices and the mortgage interest rate adjust. The increased mean consumption among these groups is mainly a result of households changing their tenure status, which also affects the composition of households in the groups defined by tenure status. The share of buyers and the homeownership rate decrease slightly in the period of the shock when allowing for both house prices and the mortgage interest rate to adjust; whereas the opposite occurs when these prices are held constant (see Figure 2.12). When house prices increase in the period of the shock, some households postpone their house purchase. Households who buy a house tend to consume less in the period of the purchase, due to transaction costs. When fewer households buy, the mean consumption response among buyers increases. Moreover, as the households who postpone a house purchase no longer need to pay the transaction costs of buying, they can increase their consumption more in the period of the shock, thus contributing to the larger mean consumption response among renters.

What channels of monetary transmission can explain the immediate increase in the house price? The marginal buyers in the model are affected by both the cash-flow channel and the PTI channel, in addition to the intertemporal-substitution channel. By purchasing a house, the cash-flow channel immediately transforms the negative effect of a low return on savings into a benefit of lower interest payments on debt. This effect makes the marginal buyer value ownership higher. In addition, the relaxed PTI requirement, which follows from the lower mortgage interest rate, can enable the marginal buyer to take up a bigger mortgage or to purchase a larger home. However, the average size of owned housing and the homeownership rate remain fairly constant in the first period of the transition, suggesting a muted effect due to a relaxed PTI constraint. These findings combined with the initial increase in house prices suggest that the marginal buyers increase their demand for owned housing due to the cash-flow channel, but they are restricted in terms of how large a house they can afford, due to the down-payment requirement.

Next, let us consider how the savings behavior is affected by the real interest rate shock. The relevant savings rate for mortgagors is their mortgage interest rate, since it is higher than the return on risk-free bonds. The main reason why mortgagors also save in bonds is for liquidity purposes. The difference between the mortgage interest rate and the return on risk-free bonds impacts the choice of how much to save in mortgages versus bonds. Immediately after the interest rate shock occurs, the deviation between the two rates is the largest, as mortgage interest rates are long-term and bond rates adjust periodically. Hence, it is favorable to allocate more savings to mortgages and less to bonds. This is exactly what we see in Figure 2.6. In addition, existing mortgagors who do not refinance their mortgage are not at all affected by the change in the mortgage interest rate. By paying off their mortgage, they save at the rate previously specified in their mortgage contract. For them, there is an even larger difference in returns between saving in liquid bonds or through their mortgage, contributing further to the shift in savings from bonds to mortgages.

Gradual effects

Although the focus of this paper is on the direct demand effects of a real interest rate shock, let us still consider how the direct effects play out over time. The channels through which the transmission of monetary policy works affect households' behavior over an extended period of time. As previously discussed, the sharp drop in consumption that follows the initial increase can mainly be explained by intertemporal substitution, and the decrease is likely larger due to that labor income does not adjust endogenously in the model.

Similarly, intertemporal substitution explains the distinct drop in house prices after the initial increase. In the period after the interest rate shock occurs, the cash-flow benefits for the marginal buyers are gone, as the mortgage interest rate has returned to the steady-state level. The fact that house prices are lower than in steady state, in the period after the shock, comes from the reduced overall resources available to the marginal buyers. Most of the marginal buyers were renters in the previous period, where they substantially reduced their savings as a result of intertemporal substitution. Further, the lower return on the liquid savings contributes additionally to less available resources for the marginal buyers.

The smaller savings of renters are also reflected in a lower homeownership rate, following the interest rate shock. The gradual decrease in homeownership is shown in Figure 2.7, along with the corresponding change in the average size of owned housing units. The small gradual increase in the average size of owned housing mainly reflects a selection effect from the lower homeownership rate. Young households who postpone their house purchase are relatively poor as compared to the average homeowner and choose smaller houses than the average when they buy.

To summarize this, the mortgage cash-flow channel and the general-equilibrium channel on house prices contribute to a small amplification in the aggregate response of consumption to a monetary policy shock. The immediate aggregate response of consumption to a real interest rate shock of -80bp is an increase by 0.09 percent, when the mortgage contract is a 30-year fixed-payment mortgage. This is 50 percent higher as compared to when house prices and mortgage interest rates are held constant. The amplification is mostly driven by liquidity-constrained homeowners who refinance their mortgage and cash-out housing equity

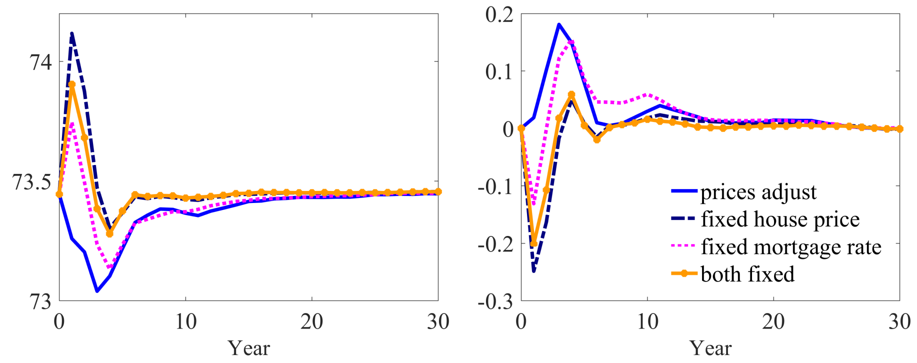

(a) Homeownership rate (%) **(b)** Mean owned housing size (%)

Figure 2.7: Impulse response functions for the homeownership rate and the average size of owned housing

Note: A decomposition of aggregate responses under different equilibrium assumptions for house prices and the mortgage interest rate. The impulse response functions follow an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rate, as displayed in Figure 2.3.

to smooth consumption. Neither changes in house prices nor mortgage interest rates alone contribute significantly to the increase, but it is the interaction that matters. Furthermore, as the mortgage rates are long-term, and fixed for current mortgagors who do not refinance, it is relatively more favorable to save in mortgages when the short-term interest rate on liquid savings declines. Therefore, there is a shift in savings from liquid bonds to mortgages, where the mortgage balance in the aggregate goes down. Young renters save less due to the lower interest rate on bonds and cause the house price and the homeownership rate to decline in subsequent periods.

2.4.2 Adjustable-rate mortgages and cash-out refinancing

Motivated by the growing empirical literature which suggests that ARMs can contribute to stronger responses to monetary policy, I now compare the IRFs in the previous section to those when ARMs are used instead of FRMs.[18] The findings from the analysis with FRMs point towards a stronger response in consumption due to all the considered transmission

[18]See, for example, Calza et al. (2013), Di Maggio et al. (2017), and Flodén et al. (2019).

channels: the intertemporal-substitution channel, the cash-flow channel, the PTI channel, and the house-price channel, when mortgage contracts have adjustable rates. There are two main differences between ARMs and FRMs. First, all mortgagors are affected by changes in the mortgage interest rate under ARMs, as opposed to only those who take up a new mortgage or refinance their mortgage under FRMs. Second, under ARMs the mortgage rates are short term. Hence, a temporary shock to the real interest rate directly translates into a change in the mortgage interest rate of the same magnitude. With FRMs, on the other hand, the mortgage interest rates are long term, and thus a temporary change in the short-term real interest rate leads to a smaller change in the long-term mortgage interest rate, but this rate applies for 30 years. The resulting mortgage interest rate path for ARMs following the same unexpected, expansionary shock to the real interest rate as in the previous section, is displayed in Figure 2.8. The figure also presents a comparison with the mortgage interest rate path under FRMs. Beyond the fixing period length of interest rates, the mortgage contracts in the two settings are equal, i.e., mortgages are amortized over 30 years, and the same LTV and PTI constraints and refinancing costs apply.

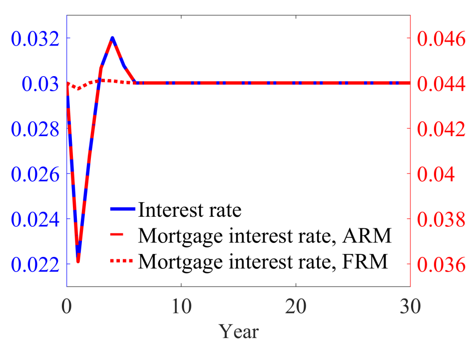

Figure 2.8: The mortgage interest rate paths
Note: A comparison of mortgage interest rates in the baseline model with fixed-rate mortgages and the comparable model with adjustable-rate mortgages. The mortgage interest rates read off the right-hand side y-axis. The paths follow an unexpected nominal interest rate shock of -100bp, where the path of the real interest rate in the figure corresponds to the estimated impulse response function in Auclert et al. (2020).

The IRFs for consumption under the two different mortgage specifications are displayed in Figure 2.9.[19] Once more, to explore the role

[19]For visual purposes, I do not show the full transition paths in the figures, but the

of endogenous changes in house prices, I also present the aggregate consumption response under ARMs when house prices are held constant. To investigate the role of refinancers, a decomposition of the immediate mean consumption response for the four groups of households (buyers, stayers, refinancers, and renters) is presented in Table 2.5.

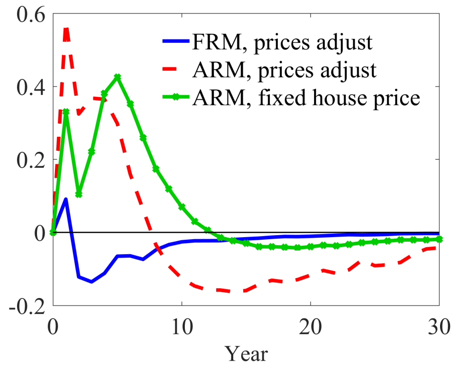

Figure 2.9: Impulse response functions for consumption (%)
Note: A decomposition of aggregate consumption responses under different equilibrium assumptions for house prices, and a comparison between the baseline model with fixed-rate mortgages and the comparable model with adjustable-rate mortgages. The impulse response functions follow an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rates, as displayed in Figure 2.8.

The immediate response of consumption is significantly larger when mortgages have adjustable rates as opposed to fixed rates. The aggregate response in consumption is an increase by 0.58 percent, which is more than six times as high as for the economy with fixed-rate mortgages. The stronger response partly stems from households who refinance their mortgage, although households who buy a house and renters also increase consumption significantly more in the economy with ARMs. Furthermore, under ARMs, the lower mortgage interest rate in itself contributes to about half of the amplification in aggregate consumption as compared to the economy with FRMs, which is seen in the last row in Table 2.5. The remaining 50 percent of the amplification comes from the interaction effect of lower mortgage interest rates and higher house prices.

Figure 2.10 presents how the share of refinancers changes in response to the interest rate shock, and well as the responses of liquid savings,

choice of the transition period length was made to ensure that all variables converge to their steady-state levels.

		Overall	Buyers	Refinancers	Stayers	Renters
FRM		0.09	0.03	0.37	0.02	0.31
	fixed p_h	0.06	-0.48	0.04	0.08	0.17
	fixed r_m	0.06	-0.18	0.09	0.05	0.18
	fixed p_h & r_m	0.06	-0.29	-0.07	0.08	0.17
ARM		0.58	1.42	5.78	-0.34	0.90
	fixed p_h	0.33	-2.69	4.48	0.22	-0.02

Table 2.5: Consumption responses in the period when the interest rate shock occurs

Note: A decomposition of mean consumption responses of buyers, refinancers, stayers, and renters, under different equilibrium assumptions for house prices and mortgage interest rates, and for different mortgage contract specifications: fixed-rate mortgages versus adjustable-rate mortgages. The deviations of consumption, in percent, are computed for the period when the real interest rate shock occurs. The separation into buyers, refinancers, stayers, and renters is based on the tenure choice in the period of the interest rate shock. The responses follow an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rates, as displayed in Figure 2.8.

mortgage balances, and the endogenous movement in house prices. Figure 2.10a shows that the refinancing rate spikes immediately after the interest rate shock occurs. Under ARMs this is solely due to a demand for cashing-out housing equity, as mortgages have variable rates. The increase in the frequency of refinancing is still motivated by the decrease in the mortgage interest rate, although the instant increase in house prices of 2.6 percent in the economy with ARMs allows for larger cash-outs by relaxing the LTV constraint. Another important observation from Figure 2.10 is that the aggregate outstanding mortgage balance increases in response to the expansionary shock. The aggregate dissaving mainly works through an increase in mortgages. In fact, the total savings in liquid bonds actually increase. This stands in sharp contrast to the model economy with FRMs, where homeowners shift savings from liquid bond to mortgages in response to the expansionary interest rate shock.

To understand the stronger consumption responses of refinancers and the higher frequency of refinancing under ARMs, let us consider the various transmission channels. In terms of the intertemporal-substitution channel and the cash-flow channel, the difference between short-term and long-term interest rates of mortgages plays an important role. Since the mortgage interest rate is the relevant savings rate for all mortgagors, including refinancers, they face a much lower savings rate under ARMs than under FRMs. It follows that mortgagors substitute consumption intertemporally

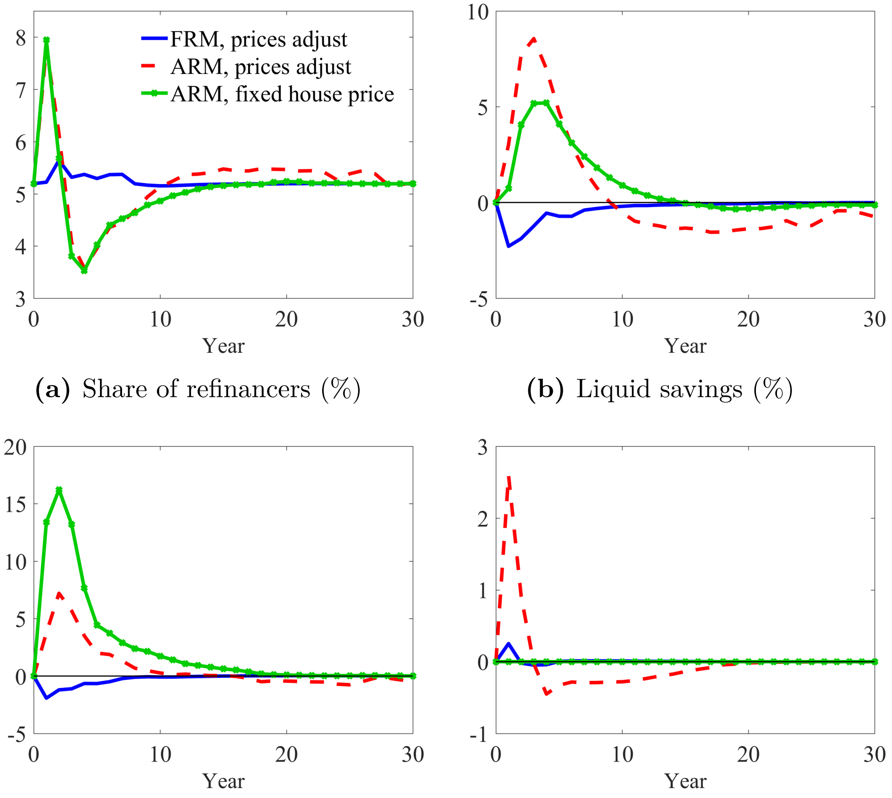

(a) Share of refinancers (%) **(b)** Liquid savings (%)

(c) Mortgage balances (%) **(d)** House prices (%)

Figure 2.10: Impulse response functions for refinancing behavior, savings in liquid bonds, mortgage balances, and house prices

Note: A decomposition of aggregate responses under different equilibrium assumptions for house prices, and a comparison between the baseline model with fixed-rate mortgages and the economy with adjustable-rate mortgages. The impulse response functions follow an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rates, as displayed in Figure 2.8.

to a greater extent when mortgages have variable rates, in response to the expansionary interest rate shock. Moreover, among the constrained households, who are not behaving according to their Euler equation, the cash-flow channel contributes to a stronger consumption response under ARMs. All mortgagors prefer the significantly lower short-term mortgage interest rates under ARMs over the more stable long-term mortgage interest rates under FRMs, since they have to pay off their mortgage over time. However, the substantially lower mortgage interest rates in the short run are particularly beneficial for the temporary constrained homeowners. Specifically, when liquidity-constrained households face much lower mortgage payments in the next-coming periods, they can reduce their precautionary savings and increase their consumption to a greater extent. However, since they are liquidity constrained, the only feasible way to reduce savings and increase consumption is through refinancing. Hence, there is both a higher frequency of refinancing as well as a larger increase in consumption among those who refinance, under ARMs. The PTI channel and the general-equilibrium channel on house prices further contribute to the stronger responses in consumption of refinancers under ARMs. The lower short-term mortgage interest rates relax the PTI constraint more than under FRMs, and the sharp increase in house prices allows for larger cash-outs.

The share of buyers in the economy increases slightly immediately after the interest rate shock, as some households want to take advantage of the temporary lower mortgage interest rate. Some of the house buyers are relatively young and use a large mortgage to finance their first house purchase, as seen in Figure 2.1. Other buyers are already homeowners who use the opportunity when the mortgage interest rate is low to buy a larger house. Mortgages with adjustable rates, where the interest rates respond more strongly in the short term to the real interest rate shock, are valued much higher by these households, compared to mortgages with fixed rates. The benefit of the much lower mortgage interest rates in the short run leads to an increase in the demand for housing, which is reflected in the higher house price. However, similarly to the economy with FRMs, the homeownership rate decreases when house prices are allowed to adjust, as seen in Figure 2.11a. This is the result of some renters postponing their house purchase due to the temporary higher house price.

The stronger consumption responses of buyers and renters when prices adjust, as seen in Table 2.5, can once again mainly be explained by households who change their tenure status (see Figure 2.14). Homeowners who decide to buy a larger house due to the lower mortgage interest rate are not very constrained in their liquidity, and therefore contribute to the stronger average consumption response among buyers. In fact, these households behave similarly to refinancers, since they can choose to take up a larger mortgage when they buy a new house. Renters who choose to postpone their house purchases, due to the temporary higher house prices, become less liquidity constrained as they no longer need to pay the transaction costs of buying. As a result, the mean consumption responses of renters and buyers are significantly higher under ARMs as opposed to FRMs, when house prices are allowed to adjust. The more muted average consumption response of stayers, when house prices adjust, partly follows from them expecting a reduction in their wealth in the future, when house prices decline.

Similarly to the economy with FRMs, intertemporal substitution and the lower return on savings lead to less available resources of renters in the years following the interest rate shock. The smaller savings of renters contribute to some households postponing their house purchases, causing the gradual decrease in homeownership. As the households who postpone their house purchases are relatively poor as compared to the average homeowner, there is a small increase in the average size of owned housing, as seen in Figure 2.11b.

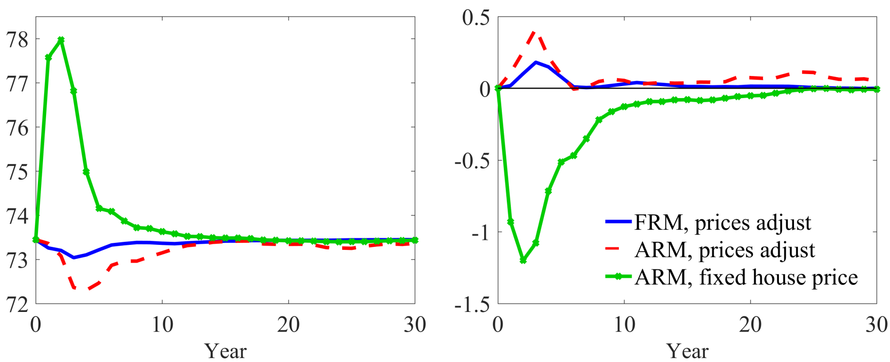

(a) Homeownership rate (%) **(b)** Mean owned housing size (%)

Figure 2.11: Impulse response functions for the homeownership rate and the average size of owned housing

Note: A decomposition of aggregate responses under different equilibrium assumptions for house prices, and a comparison between the baseline model with fixed-rate mortgages and the economy with adjustable-rate mortgages. The impulse response functions follow an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rates, as displayed in Figure 2.8.

Overall, the structure of the mortgage market substantially impacts the effectiveness of monetary policy. Adjustable-rate mortgages contribute to significantly stronger responses in consumption following an expansionary interest rate shock, as compared to when fixed-rate mortgages are used. The amplification is largely driven by constrained homeowners who refinance their mortgage and cash-out housing equity to smooth consumption. Furthermore, households who postpone purchasing a house, due to the higher house price, contribute to a larger immediate increase in consumption, since they no longer have to pay the transaction costs of buying. When mortgages have adjustable rates, the reduction in the short-term mortgage interest rate on its own contributes to about half of the amplification. The remaining 50 percent comes from the interaction of lower mortgage interest rates and higher house prices, where higher house prices enable larger cash-outs for refinancing households.

2.5 Concluding remarks

Over the past decades, there have been important developments in macroeconomic research that emphasize that different households respond very differently to changes in their environment, and that this can have implications for aggregate responses to policy changes. Many households are liquidity constrained and respond strongly to changes in their liquidity. In this paper, I explore one channel through which monetary policy can directly influence some households' cash flows, namely, through changes in their mortgage interest payments. I use a heterogeneous-agent life-cycle model to quantify the role of changes in mortgage interest rates and house prices in the transmission of monetary policy.

I find that the aggregate response of consumption to an interest rate shock is amplified due to changes in mortgage interest rates and also due to an endogenous response in house prices. Further, the structure of the mortgage market substantially impacts the effectiveness of monetary policy. Specifically, when mortgages have fixed interest rates, the direct response in aggregate consumption is about 50 percent higher due to changes in mortgage rates and house prices. However, when mortgages have adjustable rates, the aggregate response of consumption is more than six times as large as compared to when mortgages have fixed rates. The amplification is largely driven by liquidity-constrained homeowners, who expect higher earnings in the future, and who refinance their mortgage and cash-out housing equity to smooth consumption. Under adjustable-rate mortgages, the reduction in the short-term mortgage interest rate contributes to about half of the amplification. The remaining 50 percent comes from the interaction effect of lower mortgage interest rates and higher house prices, where the higher house prices allow for larger cash-outs.

The findings in this paper are in line with most of the results in recent empirical work. The results suggest that a detailed understanding of the contract structures in the mortgage market is an important input into the analysis of monetary policy. That said, a natural step to take in future work is to extend the model in this paper by including a production sector and an endogenous labor-supply choice, as well as nominal frictions. Along other dimensions, it would be interesting to investigate a possible house price path dependence of the transmission of monetary policy, and potential asymmetries in the responses to interest rate shocks with opposite signs.

References

Auclert, A. (2019, June). Monetary policy and the redistribution channel. *American Economic Review 109*(6), 2333–67.

Auclert, A., M. Rognlie, and L. Straub (2020, January). Micro jumps, macro humps: Monetary policy and business cycles in an estimated hank model. Working Paper 26647, National Bureau of Economic Research.

Bell, F. C. and M. L. Miller (2005, August). Life tables for the United States Social Security Area 1900-2100. Actuarial study No. 120, Social Security Administration.

Beraja, M., A. Fuster, E. Hurst, and J. Vavra (2018, 09). Regional heterogeneity and the refinancing channel of monetary policy. *The Quarterly Journal of Economics 134*(1), 109–183.

Boppart, T., P. Krusell, and K. Mitman (2018). Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. *Journal of Economic Dynamics and Control 89*, 68 – 92.

Calza, A., T. Monacelli, and L. Stracca (2013, 01). Housing finance and monetary policy. *Journal of the European Economic Association 11*, 101–122.

Carroll, C. D. and A. A. Samwick (1997). The nature of precautionary wealth. *Journal of Monetary Economics 40*(1), 41–71.

Chen, H., M. Michaux, and N. Roussanov (2013, September). Houses as ATMs? Mortgage refinancing and macroeconomic uncertainty. Working Paper 19421, National Bureau of Economic Research.

Cho, S.-W. and J. Francis (2011). Tax treatment of owner occupied housing and wealth inequality. *Journal of Macroeconomics 33*, 42–60.

Cloyne, J., C. Ferreira, and P. Surico (2019, 01). Monetary policy when households have debt: New evidence on the transmission mechanism. *The Review of Economic Studies*, 41–28. rdy074.

Cocco, J. F., F. J. Gomes, and P. J. Maenhout (2005). Consumption and portfolio choice over the life cycle. *Review of Financial Studies 18*(2), 491–533.

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

Di Maggio, M., A. Kermani, B. J. Keys, T. Piskorski, R. Ramcharan, A. Seru, and V. Yao (2017, November). Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging. *American Economic Review 107*(11), 3550–88.

Díaz, A. and M. J. Luengo-Prado (2008). On the user cost and homeownership. *Review of Economic Dynamics 11*(3), 584–613.

Druedahl, J. and T. H. Jørgensen (2017). A general endogenous grid method for multi-dimensional models with non-convexities and constraints. *Journal of Economic Dynamics and Control 74*, 87–107.

Eichenbaum, M., S. Rebelo, and A. Wong (2018, October). State dependent effects of monetary policy: the refinancing channel. Working Paper 25152, National Bureau of Economic Research.

Flodén, M., M. Kilström, J. Sigurdsson, and R. Vestman (2019, December). Household debt and monetary policy: Revealing the cash-flow channel. Research Paper 16-8, Swedish House of Finance.

Floetotto, M., M. Kirker, and J. Stroebel (2016). Government intervention in the housing market: Who wins, who loses? *Journal of Monetary Economics 80*, 106–123.

Garriga, C., F. E. Kydland, and R. Šustek (2017, May). Mortgages and monetary policy. *The Review of Financial Studies 30*(10), 3337–3375.

Gervais, M. (2002). Housing taxation and capital accumulation. *Journal of Monetary Economics 49*(7), 1461–1489.

Gorea, D. and V. Midrigan (2017). Liquidity constraints in the U.S. housing market. *Working Paper Series, National Bureau of Economic Research.*

Greenwald, D. (2018, January). The mortgage credit channel of macroeconomic transmission. (5184-16).

Gruber, J. W. and R. F. Martin (2003). Precautionary savings and the wealth distribution with illiquid durables. *International Finance Discussion Papers 773*.

Harding, J. P., S. S. Rosenthal, and C. Sirmans (2007). Depreciation of housing capital, maintenance, and house price inflation: Estimates from a repeat sales model. *Journal of Urban Economics 61*(2), 193–217.

Harris, E. (2005, November). Effective marginal tax rates on labor income. A CBO paper, Congressional Budget Office.

Heathcote, J., K. Storesletten, and G. L. Violante (2017). Optimal tax progressivity: An analytical framework. *Quarterly Journal of Economics 132*(4), 1693–1754.

Hedlund, A., F. Karahan, K. Mitman, and S. Ozkan (2019). Monetary policy, heterogeneity and the housing channel. *Working Paper*.

Kaplan, G., K. Mitman, and G. L. Violante (2017, August). The housing boom and bust: Model meets evidence. Working Paper 23694, National Bureau of Economic Research.

Kaplan, G., B. Moll, and G. L. Violante (2018, March). Monetary policy according to hank. *American Economic Review 108*(3), 697–743.

Kaplan, G. and G. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica 82*(4), 1199–1239.

Karlman, M., K. Kinnerud, and K. Kragh-Sørensen (2020). Costly reversals of bad policies: the case of the mortgage interest deduction. *Working Paper*.

Lowry, S. (2014, August). Itemized tax deductions for individuals: Data analysis. Crs report for congress, Congressional Research Service.

Piazzesi, M. and M. Schneider (2016). Housing and macroeconomics. *Handbook of Macroeconomics 2*, 1547–1640.

Romer, C. D. and D. H. Romer (2004, September). A new measure of monetary shocks: Derivation and implications. *American Economic Review 94*(4), 1055–1084.

Sommer, K. and P. Sullivan (2018). Implications of US tax policy for house prices, rents, and homeownership. *American Economic Review 108*(2), 241–74.

The Tax Foundation (2013). U.S. federal individual income tax rates history, 1862-2013 (nominal and inflation-adjusted brackets).

Wong, A. (2019). Refinancing and the transmission of monetary policy to consumption. *Working Paper*.

2.A Equilibrium definitions

2.A.1 Stationary equilibrium

Households are heterogeneous with respect to age $j \in \mathcal{J} \equiv \{1, 2, ..., J\}$, owner-occupied housing $h \in \mathcal{H} \equiv \{0, \underline{h}, ..., \overline{h} = \overline{s}\}$, mortgage $m \in \mathcal{M} \equiv \mathbb{R}_+$, mortgage age $ma \in \mathcal{MA} \equiv \{1, 2, ..., L\}$, permanent earnings $n \in \mathcal{N} \equiv \mathbb{R}_{++}$, and cash-on-hand $x \in \mathcal{X} \equiv \mathbb{R}_{++}$. Let $\mathcal{Z} \equiv \mathcal{H} \times \mathcal{M} \times \mathcal{MA} \times \mathcal{N} \times \mathcal{X}$ be the non-deterministic state space with $\mathbf{z} \equiv (h, m, ma, n, x)$ denoting the vector of individual states. Let $\mathbf{B}(\mathbb{R}_{++})$ and $\mathbf{B}(\mathbb{R}_+)$ be the Borel σ-algebras on $\mathbb{R}_{++}$ and $\mathbb{R}_+$, respectively, $P(\mathcal{H})$ the power set of $\mathcal{H}$, and $P(\mathcal{MA})$ the power set of $\mathcal{MA}$, and define $\mathscr{B}(\mathcal{Z}) \equiv P(\mathcal{H}) \times \mathbf{B}(\mathbb{R}_+) \times P(\mathcal{MA}) \times \mathbf{B}(\mathbb{R}_{++}) \times \mathbf{B}(\mathbb{R}_{++})$. Further, let $\mathbb{M}$ be the set of all finite measures over the measurable space $(\mathcal{Z}, \mathscr{B}(\mathcal{Z}))$. Then, $\Phi_j(Z) \in \mathbb{M}$ is a probability measure defined on subsets $Z \in \mathscr{B}(\mathcal{Z})$ that describes the distribution of individual states across agents of age $j \in \mathcal{J}$. Finally, denote the time-invariant fraction of the population of age $j \in \mathcal{J}$ by Π_j.

Definition 1. A stationary recursive competitive equilibrium is a collection of value functions $V_j(\mathbf{z})$ with associated policy functions $\{c_j(\mathbf{z}), s_j(\mathbf{z}), h'_j(\mathbf{z}), m'_j(\mathbf{z}), b'_j(\mathbf{z})\}$ for all j; prices (p_h, p_r); quantities of the total housing stock $\bar{H}$ and the total rental housing stock $\bar{S}$; and a distribution of agents' states Φ_j for all j such that:

1. Given prices (p_h, p_r), $V_j(\mathbf{z})$ solves the Bellman equation (2.14) with the corresponding set of policy functions $\{c_j(\mathbf{z}), s_j(\mathbf{z}), h'_j(\mathbf{z}), m'_j(\mathbf{z}), b'_j(\mathbf{z})\}$ for all j.

2. Given $p_h = p'_h$, the rental price per unit of housing services p_r is given by equation (2.15).

3. The quantity of the total housing stock is given by the total demand for housing services[20]

$$\bar{H} = \sum_{\mathcal{J}} \Pi_j \int_Z s_j(\mathbf{z}) d\Phi_j(Z).$$

4. The quantity of the total rental housing stock is given by the total

[20] I assume a perfectly elastic supply of both owner-occupied housing and rental units in steady state. This implies that supply always equals demand, and markets clear.

demand for rental housing services

$$\bar{S} = \bar{H} - \sum_{\mathcal{J}} \Pi_j \int_Z h_j(\mathbf{z}) d\Phi_j(Z).$$

5. The distribution of states Φ_j is given by the following law of motion for all $j < J$

$$\Phi_{j+1}(\mathcal{Z}) = \int_Z Q_j(\mathbf{z}, \mathcal{Z}) d\Phi_j(Z),$$

where $Q_j : \mathcal{Z} \times \mathscr{B}(\mathcal{Z}) \rightarrow [0, 1]$ is a transition function that defines the probability that a household at age j transits from its current state $\mathbf{z}$ to the set $\mathcal{Z}$ at age $j + 1$.

2.A.2 Transitional equilibrium

Let $\Phi_{tr,jt}(Z_t) \in \mathbb{M}$ be a probability measure defined on subsets $Z_t \in \mathscr{B}(\mathcal{Z})$ that describes the distribution of individual states across agents of age $j \in \mathcal{J}$ at time period t.

Definition 2. Given a sequence of interest rates $\{r_t\}_{t=1}^{t=\infty}$ and initial conditions $\Phi_{tr,j1}(Z_1)$ for all j, a transitional recursive competitive equilibrium is a sequence of value functions $\{V_{jt}(\mathbf{z})\}_{t=1}^{t=\infty}$ with associated policy functions $\{c_{jt}(\mathbf{z}), s_{jt}(\mathbf{z}), h'_{jt}(\mathbf{z}), m'_{jt}(\mathbf{z}), b'_{jt}(\mathbf{z})\}_{t=1}^{t=\infty}$ for all j; a sequence of prices $\{(p_{h,t}, p_{r,t})\}_{t=1}^{t=\infty}$; sequences of quantities of total housing demand $\{H_t\}_{t=1}^{t=\infty}$, total rental housing demand $\{S_t^D\}_{t=1}^{t=\infty}$, and total rental housing supply $\{S_t^S\}_{t=1}^{t=\infty}$; and a sequence of distributions of agents' states $\{\Phi_{tr,jt}\}_{t=1}^{t=\infty}$ for all j such that:

1. Given prices $(p_{h,t}, p_{r,t})$, $V_{jt}(\mathbf{z})$ solves the Bellman equation with the corresponding set of policy functions $\{c_{jt}(\mathbf{z}), s_{jt}(\mathbf{z}), h'_{jt}(\mathbf{z}), m'_{jt}(\mathbf{z}), b'_{jt}(\mathbf{z})\}$ for all j and t.

2. Given $p_{h,t}$, $p_{h,t+1}$, S_t^S, and $\bar{S}$, the rental price per unit of housing service is $p_{r,t}$ for all t, and is given by equation (2.16), where for a given t, $S = S_t^S$.

3. The housing market clears:

$$
\begin{aligned}
& H_t = \bar{H} \quad \forall t \\
\text{where} \quad & H_t = \sum_{\mathcal{J}} \Pi_j \int_{Z_t} s_{jt}(\mathbf{z}) d\Phi_{tr,jt}(Z_t) \quad \forall t \\
\text{and} \quad & \bar{H} \text{ is the total housing stock in steady state.}
\end{aligned}
$$

4. The rental market clears:

$$
\begin{aligned}
& S_t^D = S_t^S \quad \forall t \\
\text{where} \quad & S_t^D = H_t - \sum_{\mathcal{J}} \Pi_j \int_{Z_t} h'_{jt}(\mathbf{z}) d\Phi_{tr,jt}(Z_t) \quad \forall t \\
\text{and} \quad & S_t^S \text{ is the total rental housing supply in period } t.
\end{aligned}
$$

5. Distributions of states $\Phi_{tr,jt}$ are given by the following law of motion for all $j < J$ and t:

$$
\Phi_{tr,j+1,t+1}(\mathcal{Z}) = \int_{Z_t} Q_{tr,jt}(\mathbf{z}, \mathcal{Z}) d\Phi_{tr,jt}(Z_t),
$$

where $Q_{tr,jt} : \mathcal{Z} \times \mathscr{B}(\mathcal{Z}) \to [0, 1]$ is a transition function that defines the probability that a household of age j at time t transits from its current state $\mathbf{z}$ to the set $\mathcal{Z}$ at age $j + 1$ and time $t + 1$.

2.B Computational method and solution algorithm

See Karlman et al. (2020) for a detailed description of the computational method. To summarize, I use the general generalization of the endogenous grid method G^{2}EGM by Druedahl and Jørgensen (2017) to solve for the value and policy functions. The state space is discretized, where the number of grid points for permanent earnings N_N, cash-on-hand N_X, housing sizes N_H, bonds-over-earnings N_B, and loan-to-value N_{LTV}, are 9, 39, 4, 20, and 21, respectively. At lower levels of cash-on-hand and bonds-over-earnings, the grid points are denser.

All monetary policy shocks are unexpected and I adjust individual states for changes in the house price and taxes. Specifically, cash-on-hand x needs to be adjusted because (i) the value of the house changes; (ii) the property tax payment is affected; and, (iii) of changes in tax deductions due to changes in property taxes. In addition, I need to adjust for changes in the loan-to-value due to changes in the house price.

2.B.1 Solution algorithm

Steady state

Solving the steady state:

1. Impose house price $p_h = 2.60$ and compute p_r from equation (2.15).
2. Solve the household problem recursively, and obtain the value and policy functions.
3. Simulate using optimal decision rules.
4. Use simulated values to compute the total housing stock $\bar{H}$ and the total rental stock $\bar{S}$. From the simulation I also get the distribution of agents' states Φ_j for all j.

Transition

Let $\Phi_{init,j}$ be the distribution of households' states in the initial steady state. Further, let t denote the transition period, and assume that the economy has returned to steady state in $t = T+1$. Choose T large enough so that by increasing T the transition path is unaffected.[21]

Solving the transition:

1. Guess $\{p_{h,t}\}_{t=1}^{t=T}$ and $\{S_t^S\}_{t=1}^{t=T}$, and compute $\{p_{r,t}\}_{t=1}^{t=T}$ using the steady-state rental housing stock $\bar{S}$.
2. Recursively solve for the value and policy functions for all ages $j \in \mathcal{J}$ and time periods $t \in \mathcal{T}$. For $t = T+1$, take the value and policy functions from the steady state.

[21] I set $T = 30$, and $T = 40$, depending on the experiment.

3. Given the price $p_{h,1}$, for each $j \in \mathcal{J}$, adjust the initial individual states such that the initial distribution $\Phi_{init,j}$ reflects unexpected changes in the house price from the initial steady state.

4. Simulate using the adjusted initial distribution and optimal decision rules. Use simulated values to compute the sequence of total housing demand $\{H_t\}_{t=1}^{t=T}$ and total rental housing demand $\{S_t^D\}_{t=1}^{t=T}$.

5. Compute the sequence of excess demand for housing $\{ED_{H,t}\}_{t=1}^{t=T}$ and for rental housing $\{ED_{S,t}\}_{t=1}^{t=T}$, and the Euclidean norms of these sequences.

   (a) If the norm is larger than some tolerance level, update $\{p_{h,t}\}_{t=1}^{t=T}$ using the rule $p'_{h,t} = p_{h,t} + ED_{H_t} * \epsilon_{p_h}$ and $\{S_t^S\}_{t=1}^{t=T}$ using the rule $S_t^{S'} = S_t^S + ED_{S,t} * \epsilon_S$, for all $t \in \mathcal{T}$ and go back to step 1.

   (b) If the norms are within the tolerance level, convergence is achieved.

2.C Impulse response functions

2.C.1 Fixed-rate mortgages

Figure 2.12 shows how the shares of homeowners, refinancers, buyers, and stayers change in response to the interest rate shock, for different equilibrium assumptions. Figure 2.13 displays the path of the rental rate, and the average size of rental housing, following the interest rate shock, in the economy where mortgage contracts have fixed rates.

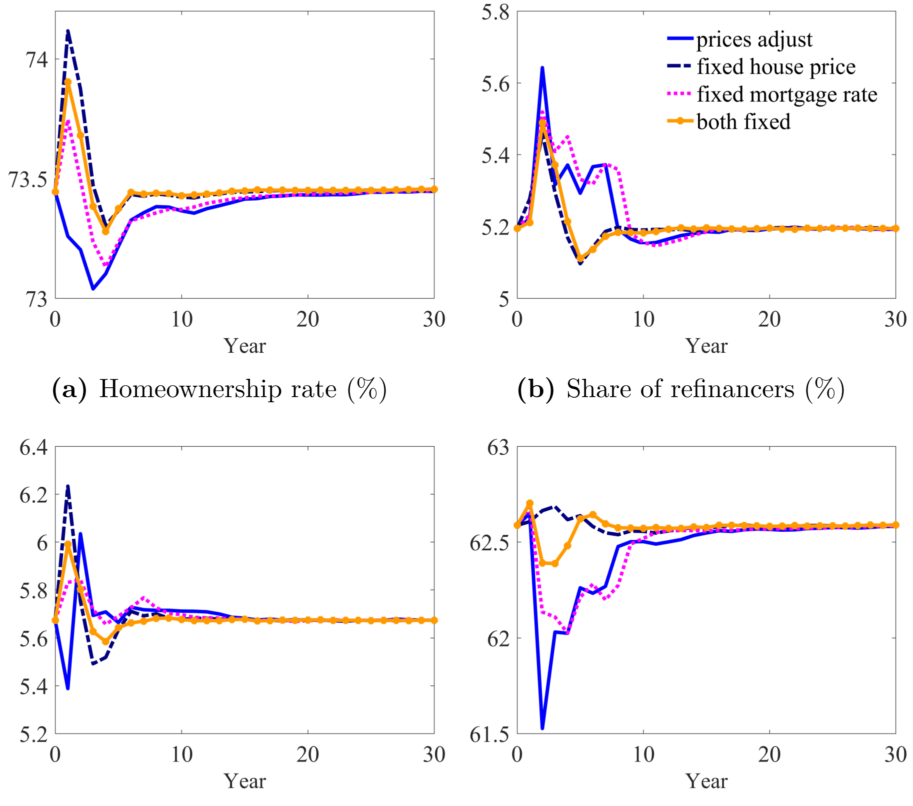

(a) Homeownership rate (%) **(b)** Share of refinancers (%)

(c) Share of buyers (%) **(d)** Share of stayers (%)

Figure 2.12: Impulse response functions for tenure statuses, under FRMs
Note: A decomposition of aggregate responses under different equilibrium assumptions for house prices and the mortgage interest rate. The impulse response functions follow an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rate, as displayed in Figure 2.3.

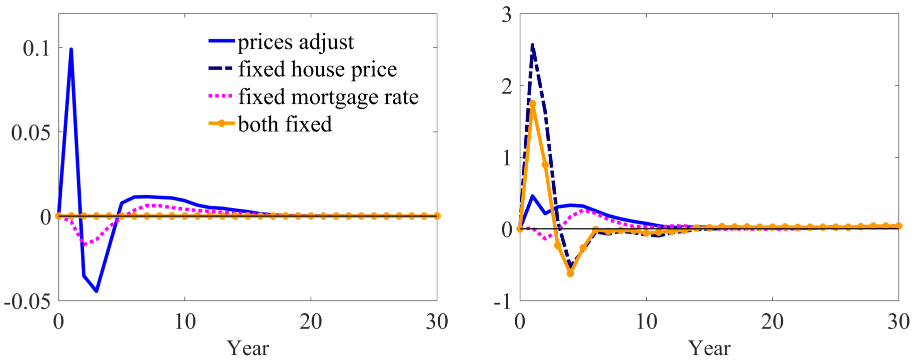

(a) The path of the rental rate (%) **(b)** Housing services, renters (%)

Figure 2.13: The path of the rental rate and the IRFs for rental services, under FRMs

Note: The rental-rate paths and the impulse response functions for the average rental size under different equilibrium assumptions for house prices and the mortgage interest rate. The impulse response functions follow an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rate, as displayed in Figure 2.3.

2.C.2 Adjustable-rate mortgages

Figure 2.14 shows how the shares of homeowners, refinancers, buyers, and stayers change in response to the interest rate shock, for different equilibrium assumptions. Figure 2.15 displays the path of the rental rate, and the average size of rental housing, following the interest rate shock, in the economy where mortgage contracts have variable rates.

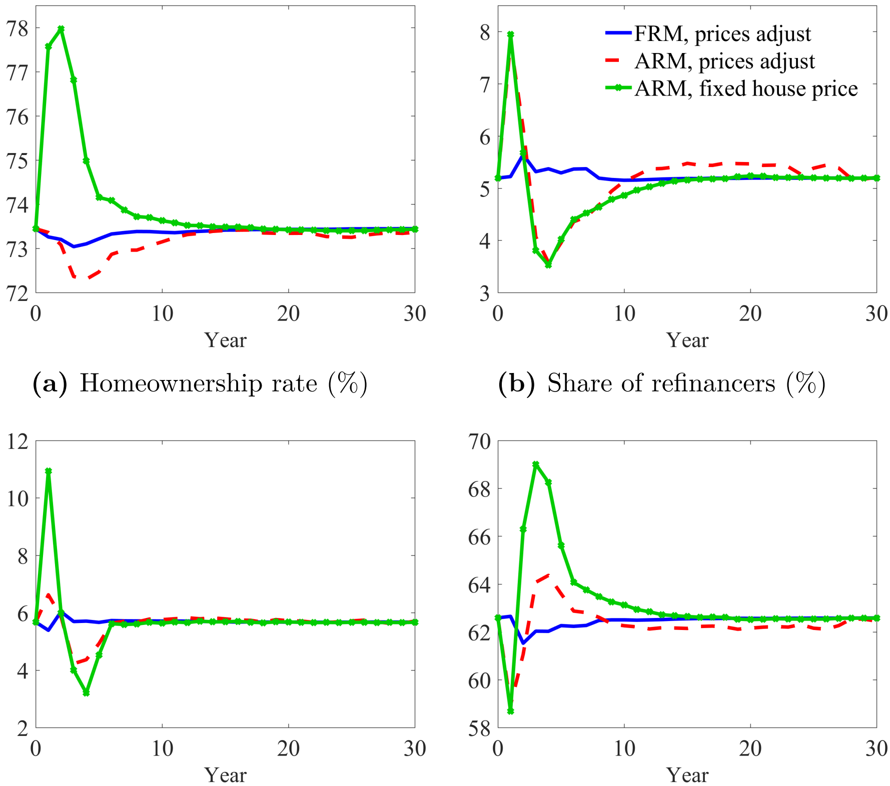

(a) Homeownership rate (%) **(b)** Share of refinancers (%)

(c) Share of buyers (%) **(d)** Share of stayers (%)

Figure 2.14: Impulse response functions for tenure statuses

Note: A decomposition of aggregate responses under different equilibrium assumptions for house prices, and a comparison between the baseline model with fixed-rate mortgages and the economy with adjustable-rate mortgages. The impulse response functions follow an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rates, as displayed in Figure 2.8.

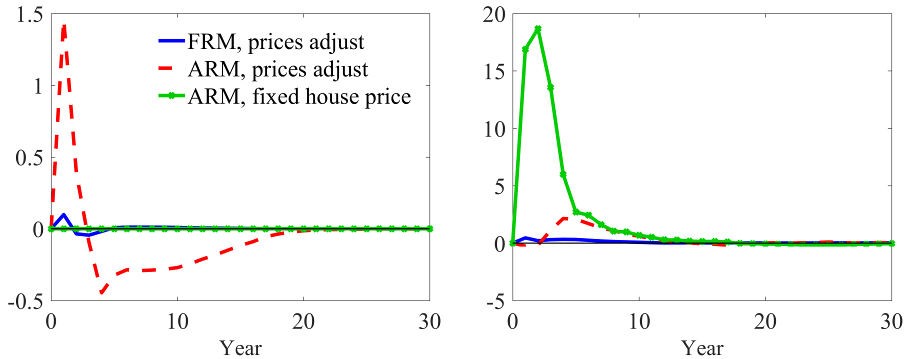

(a) The path of the rental rate (%) **(b)** Housing services, renters (%)

Figure 2.15: The path of the rental rate and the IRFs for rental services
Note: The rental-rate paths and the impulse response functions for the average rental size under different equilibrium assumptions for house prices. A comparison between the baseline model with fixed-rate mortgages and the economy with adjustable-rate mortgages. The impulse response functions follow an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rates, as displayed in Figure 2.8.

2.D Robustness: response in aggregate income

This section presents the main findings when including a response in aggregate income to the real interest rate shock. I use an empirically estimated path of output from a shock of -100bp to the nominal interest rate. The estimated path of aggregate output is the impulse response function from the identified Romer and Romer (2004) monetary policy shock in Auclert et al. (2020), which is consistent with the estimated path of the real interest rate that is used in the analysis. The resulting response in output is displayed in Figure 2.16.

There are two major caveats of this analysis. First, including a response in aggregate income that affects all households in the same way, clearly does not account for how general equilibrium effects are likely to influence different households very differently. Second, since the impulse response function for aggregate income is estimated using U.S. data, the responses are consistent with a mortgage market where the contracts are mostly of the fixed-rate type. Given the results in this paper, the direct demand effects from an interest rate shock are much larger when ARMs are used. Hence, the aggregate response in income would likely be larger in such an economy.

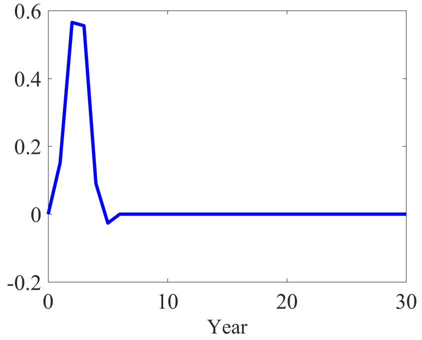

Figure 2.16: Response in aggregate income (%)
Note: The response in aggregate income follows an unexpected nominal interest rate shock of -100bp, where the path of aggregate income in the figure corresponds to the estimated impulse response function for output in Auclert et al. (2020).

The IRFs for consumption to the real interest rate shock, under the two different mortgage specifications, and when aggregate income responds as in Figure 2.16, are presented in Figure 2.17. A decomposition of the consumption responses across tenure types is provided in Table 2.6. Figure 2.18 shows how the share of refinancers changes in response to the real interest rate shock, along with the impulse response functions for savings, mortgages, and house prices; while Figure 2.19 presents the paths of the homeownership rate and the average owned housing size.

The main findings from the analysis without a response in aggregate income remain when aggregate income adjusts to the interest rate shock. That is, the immediate aggregate response in consumption is significantly larger when mortgages have adjustable rates as opposed to fixed rates. The increase in demand is largely driven by households who refinance their mortgage, but also by renters and households who buy a home in the period of the interest rate shock. Moreover, aggregate savings in liquid bonds decrease in response to the expansionary shock under FRMs, but increase when contracts have variable rates. This is accompanied by a decrease in the aggregate mortgage balance under FRMs, and an increase under ARMs.

As expected, the increase in consumption is larger under both FRMs and ARMs, when the aggregate income level rises following the interest rate shock. Furthermore, the house price level also responds stronger when aggregate income adjusts, resulting in somewhat larger changes in the homeownership rate and the mean owned housing size.

Why is it the case that the main mechanisms survive when including a response in aggregate income? As seen in Figure 2.16, the estimated increase in aggregate income arises over time. The main drivers of the amplification in aggregate consumption following the interest rate shock, are constrained homeowners, who expect higher earnings in the future. As most of the increase in earnings occurs in the period after the shock, a large prevalence of refinancing in the period of the interest rate shock is still optimal.

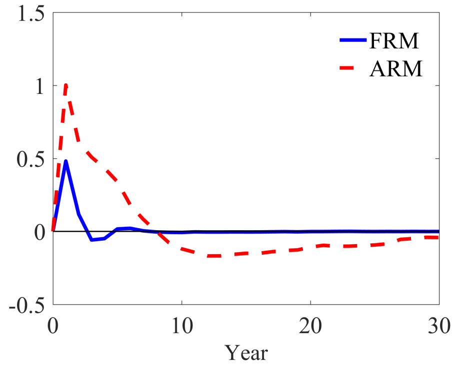

Figure 2.17: Impulse response functions for consumption (%)
Note: A comparison of aggregate consumption responses between the baseline model with fixed-rate mortgages and the comparable model with adjustable-rate mortgages. The impulse response functions follow an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rates as displayed in Figure 2.8, and changes in aggregate income as in Figure 2.16.

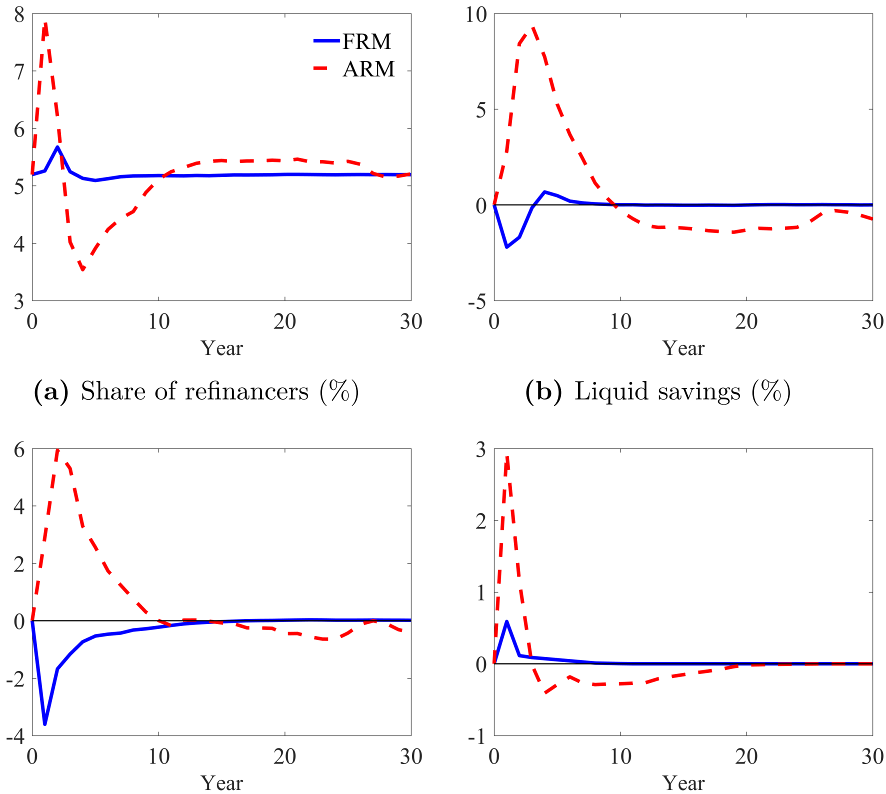

(a) Share of refinancers (%) **(b)** Liquid savings (%)

(c) Mortgage balances (%) **(d)** House prices (%)

Figure 2.18: Impulse response functions for refinancing behavior, savings in liquid bonds, mortgage balances, and house prices

Note: A comparison of aggregate responses between the baseline model with fixed-rate mortgages and the economy with adjustable-rate mortgages. The impulse response functions follow an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rates as displayed in Figure 2.8, and changes in aggregate income as in Figure 2.16.

		Overall	Buyers	Refinancers	Stayers	Renters
FRM						
	income adjusts	0.48	0.80	1.36	0.27	0.93
	fixed income	0.09	0.03	0.37	0.02	0.31
ARM						
	income adjusts	1.00	2.24	6.51	-0.04	1.43
	fixed income	0.58	1.42	5.78	-0.34	0.90

Table 2.6: Consumption responses in the period when the interest rate shock occurs

Note: A decomposition of mean consumption responses of buyers, refinancers, stayers, and renters, and for different mortgage contract specifications: fixed-rate mortgages versus adjustable-rate mortgages. The deviations of consumption, in percent, are computed for the period when the real interest rate shock occurs. The separation into buyers, refinancers, stayers, and renters is based on the tenure choice in the period of the interest rate shock. The responses follow an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rates as displayed in Figure 2.8. When income adjusts, the changes in aggregate income are as in Figure 2.16.

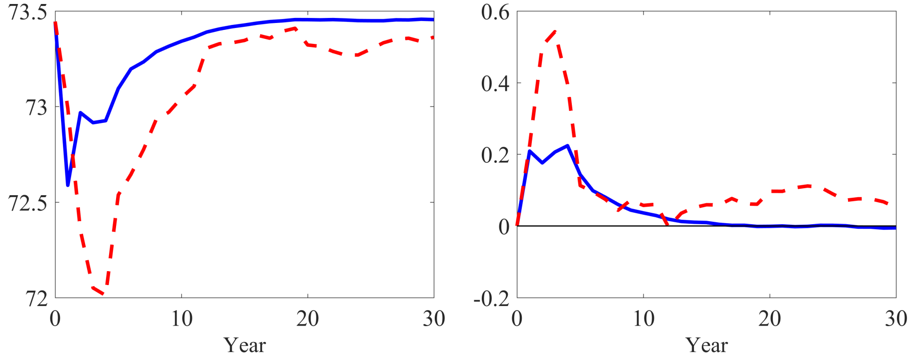

(a) Homeownership rate (%) **(b)** Mean owned housing size (%)

Figure 2.19: Impulse response functions for the homeownership rate and the average size of owned housing

Note: A comparison of aggregate responses between the baseline model with fixed-rate mortgages and the economy with adjustable-rate mortgages. The impulse response functions follow an unexpected shock to the real interest rate on bonds, with the corresponding changes in the mortgage interest rates as displayed in Figure 2.8, and changes in aggregate income as in Figure 2.16.

2.E Robustness: three-year periods

2.E.1 Calibration: model period length of three years

Parameter	Description	Value
σ^2_α	Fixed effect	0.156
σ^2_ζ	Permanent	0.035
σ^2_ν	Transitory	0.066

Table 2.7: Estimated variances
Note: The three variances are the estimated variances for: the fixed-effect earnings shock that households realize when they enter the economy, and the permanent and transitory earnings shocks to which households are subject before retirement. Estimated using PSID data.

Parameter	Description	Value	Target moment	Data	Model
α	Consumption weight	0.76	Median house value-to-earnings	2.30	2.30
β	Discount factor	0.93	Median LTV	0.35	0.35
δ^r	Depreciation rate, rentals	0.048	Homeownership rate, age $<$ 35	0.44	0.44
$\underline{h}$	Min. owned house value	0.90	Homeownership rate	0.70	0.72
ς^r	Fixed refinancing cost	0.025	Refinance rate	0.08	0.08
$\bar{q}$	Luxury of bequests	5.1	Net worth p75/p25, age 68-76	5.37	5.54
υ	Utility shifter of bequests	14.5	Median net worth, age 75/50	1.44	1.59
SD	Standard deduction	0.090	Itemization rate	0.53	0.53
λ	Level, tax function	0.975	Average marginal tax rates	0.13	0.13
τ^p	Progressivity, tax function	0.17	Distr. of marginal tax rates	See text	

Table 2.8: Estimated parameters
Note: Estimated parameters using simulated method of moments. The resulting parameter values are shown in column three. Column five displays the relevant target moment value in the data, while column six shows the comparable moment value in the model when the listed parameter values are used. The values are annual when relevant. The minimum owned house size $\underline{h}$, the fixed refinancing cost, the luxury parameter in the utility function for bequests, and the standard deduction SD, can be evaluated relative to the mean of expected periodic earnings during working life that is normalized to one.

2.E.2 Main findings: model period length of three years

An earlier version of this paper used a model period length corresponding to three years. I now provide the main findings from that analysis, and discuss which results are robust and which are not to changing the model period length. In this part of the analysis I use an unexpected reduction of the real interest rate of 100 basis points, a reduction that has a yearly persistence of 0.6.[22] The main differences between the real interest rate shock considered in this section and the estimated path of the real interest rate are that the estimated decline is less persistent and it is followed by a subsequent increase. Hence, the experiments in the two models are not directly comparable. However, the mechanisms of monetary policy transmission can still be compared. The paths of the real interest rate and the mortgage interest rates under the FRM and the ARM regimes are presented in Figure 2.20.

A comparison of the impulse response functions for aggregate consumption between the economy with FRMs and ARMs is displayed in Figure 2.21. A decomposition of the consumption responses across tenure types and for different equilibrium assumptions for house prices and mortgage interest rates is provided in Table 2.9.

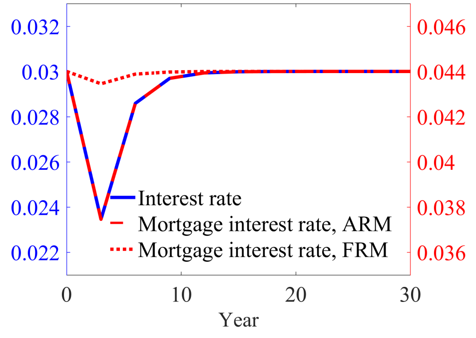

Figure 2.20: The mortgage interest rate paths
Note: A comparison of mortgage interest rates in the baseline model with fixed-rate mortgages and the comparable model with adjustable-rate mortgages. The mortgage interest rates read off the right-hand side y-axis. The paths follow an unexpected real interest rate shock of -100bp.

The main findings remain when the model period length is three years instead of one. In the baseline scenario, where households use 30-year mortgages with fixed rates, there is a small amplification in the initial consumption response due to changes in mortgage interest rates and house prices. The aggregate immediate response of consumption to

[22]The autocorrelation of the 3-month constant maturity nominal Treasury bill rates minus CPI inflation between 1982 and 2019 is 0.94 quarterly, and 0.79 annually. Garriga et al. (2017) calibrate the quarterly persistence of a real interest rate shock from the autocorrelation of the long-short spread, and find this to be 0.903, which is equivalent to an annual persistence of 0.66.

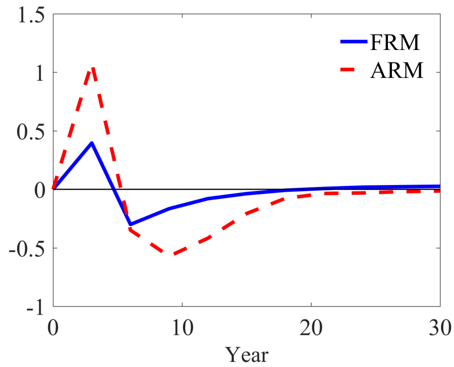

Figure 2.21: Impulse response functions for consumption (%)
Note: A comparison of the aggregate consumption responses in the baseline model with fixed-rate mortgages and the comparable model with adjustable-rate mortgages. The impulse response functions follow an unexpected real interest rate shock of -100bp.

		Overall	Buyers	Refinancers	Stayers	Renters
FRM		0.40	0.34	0.83	0.38	0.37
	fixed p_h	0.34	0.22	0.70	0.35	0.30
	fixed r_m	0.35	0.31	0.34	0.36	0.35
	fixed p_h & r_m	0.33	0.31	0.33	0.33	0.33
ARM		1.08	0.37	3.74	0.45	0.70
	fixed p_h	0.59	-0.61	3.09	0.42	0.00

Table 2.9: Consumption responses in the period when the interest rate shock occurs (%)
Note: A decomposition of mean consumption responses of buyers, refinancers, stayers, and renters, under different equilibrium assumptions for house prices and mortgage interest rates, and for different mortgage contract specifications: fixed-rate mortgages versus adjustable-rate mortgages. The deviations of consumption, in percent, are computed for the period when the real interest rate shock occurs. The separation into buyers, refinancers, stayers, and renters is based on the tenure choice in the period of the interest rate shock. The responses follow an unexpected real interest rate shock of -100bp.

the expansionary interest rate shock is approximately 20 percent higher (0.40 percent vs 0.33 percent) as a result of changes in mortgage interest rates and house prices. Further, the aggregate consumption response is substantially larger when mortgages have variable rates as opposed to fixed rates. In the economy with ARMs, the aggregate immediate response of consumption is almost three times as large, as compared to when FRMs are used (1.08 percent vs 0.40 percent).

As suggested by Table 2.9, the mechanism behind the amplification in aggregate demand is similar to that in the yearly model: the amplification is driven by the behavior of refinancers. Homeowners who refinance their mortgage increase consumption instantly by 0.83 percent under FRMs and by 3.74 percent under ARMs. Figure 2.22 shows how the share of refinancers changes in response to the real interest rate shock, and the impulse response functions for savings, mortgages, and house prices are presented. Once more, the share of refinancers increases following the decline in the real interest rate, and much more so when mortgages have adjustable rates. This follows from a substantially larger immediate increase in house prices under ARMs of approximately 3.0 percent, as compared to 0.1 percent under FRMs. The savings patterns in response to the expansionary shock are also similar to those in the previous analysis. When FRMs are used, homeowners save less in bonds and instead save by paying off more on their mortgage. With ARMs, on the other hand, savings in bonds increase following the shock; however, this is accompanied by a substantial increase in mortgage balances.

Figure 2.23 presents the paths of the homeownership rate and the average owned housing size. With the yearly model there is a gradual decrease in the homeownership rate, accompanied by a small increase in the average size of owned housing, following the expansionary real interest rate shock. However, when the model period length is three years, the pattern is the opposite. Further, there is a much slower convergence back to steady state when the period length is three years.

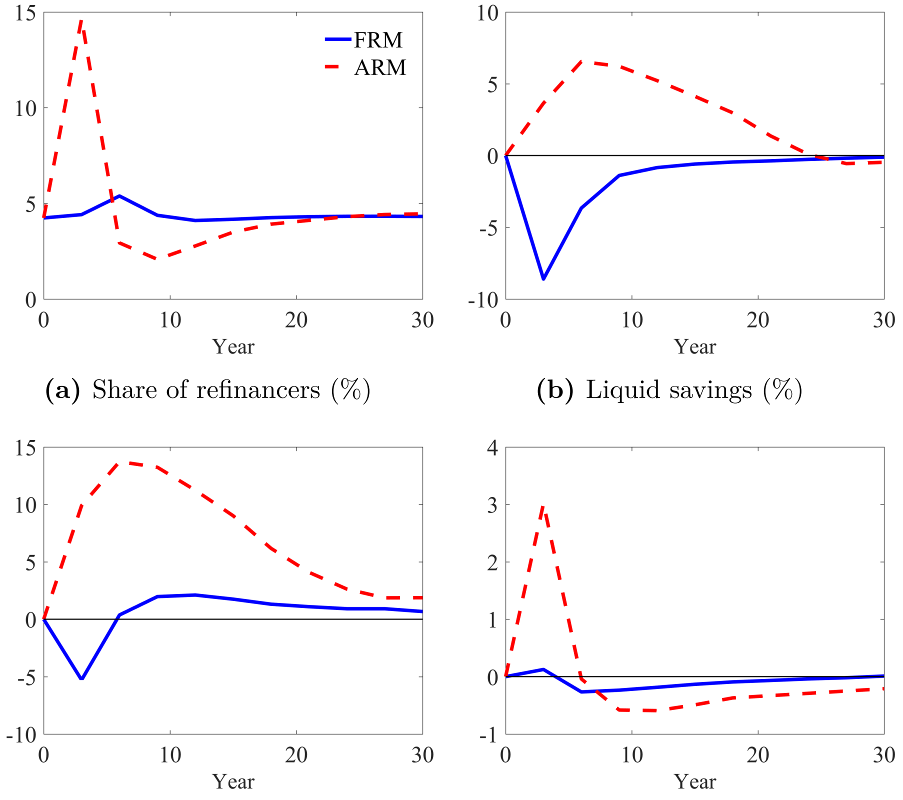

(a) Share of refinancers (%)
(b) Liquid savings (%)
(c) Mortgage balances (%)
(d) House prices (%)

Figure 2.22: Impulse response functions for refinancing behavior, savings in liquid bonds, mortgage balances, and house prices

Note: A comparison of aggregate responses in the baseline model with fixed-rate mortgages and the economy with adjustable-rate mortgages. The impulse response functions follow an unexpected real interest rate shock of -100bp.

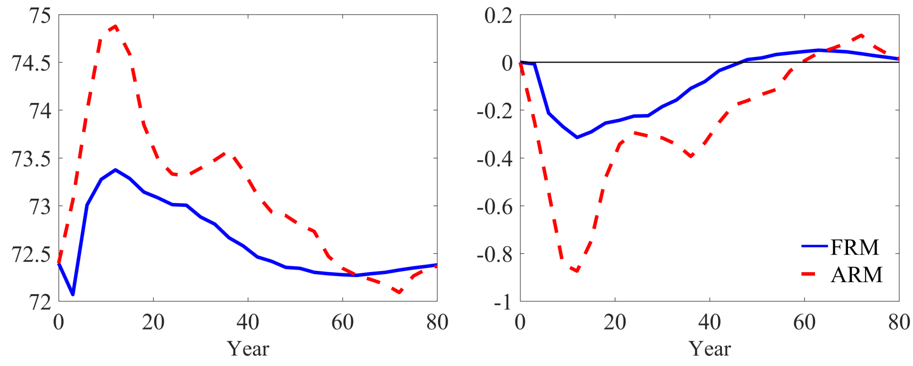

(a) Homeownership rate (%) **(b)** Mean owned housing size (%)

Figure 2.23: Impulse response functions for the homeownership rate and the average size of owned housing

Note: A comparison of aggregate responses in the baseline model with fixed-rate mortgages and the economy with adjustable-rate mortgages. The impulse response functions follow an unexpected real interest rate shock of -100bp.

Chapter 3

Mortgage lending standards: implications for consumption dynamics*

*This paper has been jointly written with Markus Karlman and Kasper Kragh-Sørensen. We are thankful for helpful discussions with John Hassler, Per Krusell, Kurt Mitman, Monika Piazzesi, Martin Schneider, Karl Walentin, and seminar participants at Stockholm University. We gratefully acknowledge funding from Handelsbanken's Research Foundations, Stiftelsen Carl Mannerfelts Fond, and Torsten Söderbergs Stiftelse. All errors are our own.

3.1 Introduction

Since the Great Recession, there has been an increased concern that high household debt makes the economy more vulnerable to adverse events. This concern partly stems from findings in the literature on the causes of the recession.[1] A prominent result in this line of work is that the rise in household debt in the early 2000's led to a stronger consumption response among households when the crisis hit. Policymakers in many countries have reacted to these findings by introducing stricter lending regulations, with the ambition to reduce the sensitivity of consumption to future shocks. As mortgages are the most common type of debt contract held by households, they have received special attention.[2]

It is not obvious, however, that stricter mortgage regulations dampen the consumption responses. First, by constraining how much households can borrow, households may find it *more* difficult to smooth consumption as their access to credit is reduced. Second, a household that chooses to take up less debt due to new regulations may also respond by lowering its buffer of liquid savings. Thus, households may adjust their asset holdings such that they are no better prepared to handle unexpected shocks.

In this paper, we study whether stricter mortgage lending standards affect consumption responses to shocks. Specifically, we investigate to what extent a permanent or temporary tightening of loan-to-value (LTV) and payment-to-income (PTI) requirements influences households' marginal propensity to consume (MPC) out of a wealth shock.

We have two main findings. First, we show that permanent policies do not materially affect aggregate consumption dynamics. In fact, a permanent tightening of the LTV or PTI constraint only marginally affects the distribution of MPCs across households. Second, a temporary one-period policy, implemented in a year prior to a negative wealth shock, can successfully reduce the consumption fall during the bust. However, such policies are, on average, only beneficial to households under very particular circumstances. The negative wealth shock needs to be large, and the policymaker must have an informational advantage in that she

[1] There is a rich literature that studies the causes of the Great Recession and the role of relaxed lending standards, through, for example, securitization of mortgage debt, and increased household debt. See, for example, Mian and Sufi (2014).

[2] For example, Sweden has implemented stricter guidelines on loan-to-income.

can perfectly foresee the bust, whereas households cannot.

To explore the role of mortgage lending standards for consumption dynamics, we use a heterogeneous-household model that includes housing and long-term mortgages. Since housing tenure and mortgage choices are strongly linked to age, we explicitly model the life cycle. Further, markets are incomplete in the sense that there is idiosyncratic earnings risk that is not fully insurable. Households derive utility from non-durable consumption goods and housings services, where housing services can be obtained by either renting or owning a house. A household can save in liquid, risk-free bonds, but also in housing. Importantly, housing equity is illiquid. First, there are transaction costs associated with both buying and selling a house. Second, there are LTV and PTI constraints that limit the size of new mortgages. Finally, it is costly to use cash-out refinancing to access housing equity.

The model produces a rich distribution of marginal propensities to consume across households.[3] Portfolio choices, both in terms of leverage and liquid bond holdings, play an important role in determining households' MPC. A significant portion of renters hold no or very little liquid bonds and are so-called *poor hand-to-mouth* households with high MPCs. Moreover, a substantial fraction of homeowners have most of their wealth in illiquid housing, as the return on housing is higher than for risk-free bonds. These households resemble the *wealthy hand-to-mouth*, as described in Kaplan and Violante (2014). However, not every homeowner with low bond savings behaves as a hand-to-mouth consumer. Some homeowners expect to pay off more on their mortgage than what is stipulated by their amortization plan, and can thus choose to costlessly pay off less in response to an adverse shock. As a result, they endogenously choose to hold small amounts of liquid bonds, but are not liquidity constrained. Lastly, households who change their discrete choice, e.g., become renters instead of buying a home in response to a negative wealth shock, tend to have large and negative MPCs.

[3] We compute MPC as the change in non-durable consumption in response to an unexpected shock to wealth (cash-on-hand), relative to the size of the shock. The use of the word *marginal* is clearly abused, since we consider shock sizes of varying magnitudes, some of which are quite large. Further, to focus on the direct effects on demand, we abstract from possible propagation mechanisms through changes in, e.g., prices caused by the wealth shocks.

To quantify the effects of introducing permanently stricter lending standards, we study two considerable changes in the LTV and PTI requirements. In the LTV experiment, homeowners can only borrow up to 70 percent of the value of their home instead of the baseline limit of 90 percent. In the PTI experiment, we lower the maximum ratio of housing-related expenses to earnings that is allowed when taking up a new mortgage, from 0.28 to 0.18.[4] Both policies cause significant changes in the economy. For example, with the stricter LTV requirement, the homeownership rate falls by seven percentage points and the median LTV among homeowners is more than halved.

Despite the considerable changes in policies, we find very small changes in both the aggregate consumption response and the distribution of MPCs across households. This holds for negative wealth shocks of various magnitudes, as well as for larger changes in the lending standards. The main reason for the small differences in MPCs is that households' precautionary savings in the long run are primarily driven by the income risk to which households are exposed and by deep parameters, e.g., households' risk aversion.

In a second round of experiments, we study the effects of LTV and PTI requirements that are temporarily tightened for one period. In these experiments, the negative wealth shock materializes in the period when the constraint returns to its baseline value. A temporary policy of this kind causes some households to save more than they otherwise would, which makes them react less strongly to the bust.

Although temporary policies do affect consumption responses to wealth shocks, there is a trade-off in terms of welfare. On the one hand, households can potentially benefit as the increased savings may make them better equipped to handle a negative wealth shock. On the other hand, temporary policies restrict consumption in the year prior to the bust, and households may already save sufficiently for precautionary reasons. Thus, the temporary policies produce both winners and losers. The winners are mainly households who abstain from buying, and thereby avoid being liquidity constrained during the bust. The losers are typically households with low earnings realizations in the year prior to the bust. These house-

[4]For each of these experiments, we solve for a new steady state and the house price changes endogenously to clear the housing market.

holds want to extract housing equity through cash-out refinancing, but the possibility to do so is limited by the policies. Overall, we find that a temporary tightening of mortgage lending standards is only welfare improving on average under certain conditions. First, the negative wealth shock must be very large. Second, a policymaker needs to have an informational advantage in terms of predicting the bust.

This paper is related to the growing strand of literature highlighting how differences in liquidity across asset classes play an important role for a broad range of macroeconomic questions. In their seminal contribution, Kaplan and Violante (2014) show that the inclusion of an illiquid asset is key for producing the high MPCs among wealthy households that are observed in data. We focus our attention on one specific type of illiquid asset, housing, and construct a model with detailed housing and mortgage markets to consider changes in mortgage lending standards. Boar et al. (2020) provide a thorough analysis of the constraints in the U.S. housing market. They show that mortgage forbearance policies, which provide relief to households with a temporary low income, can be welfare improving. Consistent with their findings, we show that households in need of refinancing, i.e., those with a low transitory income, are significantly hurt by temporary stricter LTV and PTI requirements. Greenwald (2018) finds that PTI requirements are more effective than LTV limits in counteracting cyclicality, and highlights their role in the Great Recession. Our model includes a richer heterogeneity among households, which allows us to explore differences in consumption responses across households. Moreover, we consider both permanent and temporary stricter LTV and PTI limits.

On the empirical side, Lim et al. (2011) perform cross-country regressions and find that stricter LTV and debt-to-income limits are linked to a lower cyclicality of debt. Aastveit et al. (2020) show that stricter LTV limits in Norway are associated with lower debt levels, but also a fall in liquid savings, thereby having an uncertain effect on financial vulnerability. This result is much in line with our findings.

There are also a number of papers that consider macroprudential policies and their interactions with monetary policy, of which Angelini et al. (2012) provide a review. Ferrero et al. (2018) focus on the interaction between LTV requirements and monetary policy, and find that the optimal LTV limits are countercyclical. Using a model with richer heterogeneity on the household side and a more detailed mortgage market, we confirm their findings that countercyclical policies can dampen consumption fluctuations. We further emphasize that this result requires strong assumptions on the information availability of policymakers.

The remainder of the paper is organized as follows. In Section 3.2 we describe the model, followed by a calibration and comparison to the data in Section 3.3. Section 3.4 presents the results, and Section 3.5 concludes the paper.

3.2 Model

To study how changes in mortgage lending standards affect the consumption responses of households to shocks, we build a life-cycle model with heterogeneous households and incomplete markets. Households differ in terms of their age, earnings, wealth, housing tenure status, housing wealth, and mortgage debt. Importantly, housing wealth is illiquid due to transaction costs in the housing market as well as debt constraints in the mortgage market. Specifically, households face loan-to-value (LTV) and payment-to-income (PTI) constraints when taking up a new mortgage. To further capture the constraints in the U.S. housing market, mortgages are long-term and subject to amortization plans. To smooth consumption, households may use cash-out refinancing to access their housing equity, but this comes at a cost.

The assets in the model are houses and risk-free liquid bonds. The only source of debt is mortgages. The supply of both mortgages and bonds is fully elastic, and the returns are exogenous. The aggregate housing supply, on the other hand, is inelastic and consists of both owned and rental housing units that are available in discrete sizes. In steady state, the house and rental prices adjust to clear the housing market. In addition to households, there are rental firms that provide rental housing services, and there is a government that taxes the agents and provides social security. Time is discrete, and a model period corresponds to one year. Overall, the model shares many features with the model in Karlman et al. (2020).

3.2.1 Households

The model is a life-cycle model with overlapping generations. There is a unit measure of households i of each age j. When households enter the economy at age $j = 1$, they are provided with different levels of initial net worth. The distribution of net worth among the entering cohort is matched to data, as in Kaplan and Violante (2014). Throughout their lives, households are subject to idiosyncratic earnings risk, consisting of permanent and transitory shocks. There are also age-dependent and households-specific fixed components of earnings. At age J_{ret}, households retire, and from then on they receive social security benefits that are only a share of their permanent earnings in the period before retirement, subject to a cap. In retirement, there is no permanent earnings uncertainty, but households still face transitory income shocks to proxy for expenditure shocks that older people often experience. Households in retirement face an age-dependent probability of surviving to the next period $\phi_j \in [0, 1]$, where $\phi_J = 0$.

In each period, households choose how much to consume of non-durable consumption c and housing services s. Non-durable consumption is the numeraire good in the model. Housing services can be obtained either by renting at a unit price p_r, or by owning a house at a unit price p_h. There is a linear technology that transforms owned housing units h' to housing services s, such that $s = h'$ if $h' > 0$.[5] Thus, homeowners themselves enjoy the full housing services provided by their house and are not allowed to rent out part of their property.

Households have two ways of saving. One is to buy risk-free bonds b', the other is to invest in housing. While the housing supply is fixed in the aggregate, it is flexible in its composition of rental housing and owned housing. There is a set of discrete house sizes available for rent $S = \{\underline{s}, s_2, s_3, ..., \bar{s}\}$. The sizes available for ownership constitute a proper subset H of those available for rent. Specifically, the smallest housing size available for purchase is larger than the smallest size available for rent.[6] There are transaction costs associated with both buying and selling

[5]Primes indicate the current period choice of variables that affect next period's state variables.

[6]It is common in the literature to restrict homeownership and create a selection of wealthier households among homeowners by limiting the smallest size available for purchase; see for example Cho and Francis (2011), Floetotto et al. (2016), Gervais

a house. These costs are proportional to the house value, and are given by the parameters ς^b and ς^s, respectively.

If a household chooses to purchase a house, it can take up a long-term, non-defaultable mortgage m'. The interest rate on mortgages r^m is strictly larger than the interest rate r on bonds. A mortgage has an age-dependent repayment plan that specifies the minimum payment to be made in each period. Specifically, χ_j is the share of a mortgage that needs to be paid by a household of age j, where

$$\chi_j = \left(\sum_{k=1}^{M_j} \left[\frac{1}{(1+r_m)^k} \right] \right)^{-1}. \tag{3.1}$$

M_j denotes the maturity of the mortgage. To imitate the most commonly used mortgage contract in the U.S., the 30-year fixed-payment mortgage, the maturity is set to $M_j = \min\{30, J-j\}$. This specification stipulates that the repayment period cannot extend beyond the age of certain death, thus capturing the fact that older people tend not to take up long-term mortgages. A household that wishes to deviate from the minimum-payment schedule provided in equation (3.1) can use cash-out refinancing by paying a fixed cost ς^r.

The use of mortgage financing is further limited by LTV and PTI constraints. Whenever a household takes up a new mortgage, either when buying a new home or when using cash-out refinancing, these constraints need to be fulfilled. The LTV requirement states the maximum allowable mortgage as a fraction $1-\theta$ of the house value,

$$m' \leq (1-\theta) p_h h'. \tag{3.2}$$

The payment-to-income (PTI) constraint, on the other hand, restricts the use of a mortgage by specifying that housing-related payments, including mortgage payments, cannot exceed a share ψ of current permanent earnings z,

$$\chi_{j+1} m' + (\tau^h + \varsigma^I) p_h h' \leq \psi z. \tag{3.3}$$

The housing-related payments also include property taxes τ^h, and home

(2002), and Sommer and Sullivan (2018).

insurance payments ς^I, both proportional to the house value.[7]

Households have CRRA preferences over a Cobb-Douglas aggregator of non-durable consumption and housing services.

$$U_j(c, s) = e_j \frac{\left(c^\alpha s^{1-\alpha}\right)^{1-\sigma}}{1-\sigma}, \tag{3.4}$$

where e_j is an age-dependent utility shifter that captures the tendency of household size to vary with the life cycle (see, e.g., Kaplan et al. (2020)). Further, we include a warm-glow bequest motive for households in retirement. The utility from bequests is given by

$$U^B(q') = \upsilon \frac{(q')^{1-\sigma}}{1-\sigma} \quad \text{for } j \in [J_{ret}, J], \tag{3.5}$$

where υ controls the strength of the bequest motive, and bequests q' are given by the net worth of a household, deflated by a price index $\alpha + (1-\alpha)p_h$,

$$q' = \frac{b' + p_h h' - m'}{\alpha + (1-\alpha)p_h}. \tag{3.6}$$

By deflating, a household takes into account the purchasing power of the bequests.

There are five state variables in the household problem: age j, permanent earnings z, mortgage m, house size h, and cash-on-hand x. The state variable cash-on-hand x is defined as

$$x \equiv \begin{cases} (1+r)b - (1+r^m)m + y - \Gamma - \delta^h h + (1-\varsigma^s)p_h h & \text{if } j > 1 \\ y - \Gamma + a & \text{if } j = 1, \end{cases} \tag{3.7}$$

where y is current period earnings or social security benefits, depending on the age of the household; Γ captures all taxes paid by a household; $\delta^h h$ is a maintenance cost that a homeowner has to pay, which is modeled as proportional to the house size; $(1-\varsigma^s)p_h h$ is the value of a house net

[7] The home insurance payment is only included in the PTI requirement for calibration purposes, as it is an important cost for most homeowners, but it does not enter the budget constraint of the household.

of the transaction cost for selling the house; and finally, a represents the initial assets of the newborn cohort.

The households face three different taxes. The total tax payment Γ of a household includes social security taxes, property taxes on owned housing, and labor income taxes.

$$\Gamma \equiv \mathbb{I}^w \tau^{ss} y + \tau^h p_h h + T(\tilde{y}), \tag{3.8}$$

where the social security tax is paid only by the working age population, as indicated by the dummy variable $\mathbb{I}^w$. The labor income tax is modeled by the progressive tax and transfer function $T(\tilde{y})$, which takes taxable labor income after deductions $\tilde{y}$ as its argument. For a richer description of the tax system, see Section 3.2.3.

To solve the household problem, we compute the value function in each period separately for four mutually-exclusive discrete cases related to the housing tenure choice of the household. A household can choose to rent a house (R), buy a home (B), stay in an owned house that it enters the period with and follow the repayment plan of any outstanding mortgage (S), or stay in an owned house and take up a new mortgage by refinancing (RF). In each period, the household chooses the tenure status that yields the highest value. The renter case is characterized by a household choosing not to own a house, and it is therefore not allowed to take up a mortgage, i.e., $h' = m' = 0$. In the buyer case, the household buys a new house of a different size than the previous one, i.e., $h' > 0$ and $h' \neq h$. In the stayer and refinancing cases, a household chooses to stay in the owned house it enters the period with, i.e., $h' = h$.

For each $k \in \{R, B, S, RF\}$, the household problem is characterized by the following Bellman equation, where β is the discount factor, and the set of constraints listed below. Formally,

$$V_j^k(z, x, h, m) = \max_{c,s,h',m',b'} U_j(c, s) + \beta W_{j+1}(z', x', h', m')$$

where

$$W_{j+1}(z', x', h', m') = \begin{cases} \mathbb{E}\left[V_{j+1}(z', x', h', m')\right] & \text{if } j < J_{ret} \\ \phi_j \mathbb{E}\left[V_{j+1}(z', x', h', m')\right] + (1 - \phi_j) U^B(q') & \text{otherwise} \end{cases}$$

subject to

$$\underbrace{c + b' + \mathbb{I}^R p_r s + \mathbb{I}^B (1+\varsigma^b) p_h h' + \mathbb{I}^{RF,S}(1-\varsigma^s) p_h h + \mathbb{I}^{RF}\varsigma^r}_{\text{“Expenditures”}} \leq \underbrace{x + m'}_{\text{“Money to spend”}} \tag{3.9}$$

$$\begin{aligned}
&\mathbb{I}^{B,RF} m' \leq (1-\theta) p_h h' && \text{LTV constraint}\\
&\mathbb{I}^{B,RF}\left(\frac{\chi_{j+1} m' + (\tau^h + \tau^I) p_h h'}{z}\right) \leq \psi && \text{PTI constraint}\\
&\mathbb{I}^S m' \leq (1+r_m) m - \chi_j m && \text{Min payment}\\
&s = h' && \text{if } h' > 0\\
&m' \geq 0 && \text{if } h' > 0\\
&m' = 0 && \text{if } h' = 0\\
&c > 0, s \in S, h' \in H, b' \geq 0.
\end{aligned}$$

Equation (3.9) states the household's budget constraint. The variables $\mathbb{I}^k$ are indicator variables that equal one for the relevant tenure status case $k \in \{R, B, S, RF\}$, and zero otherwise. These capture that only renters pay rent, only refinancers pay the refinancing cost, and only if you buy or sell a house do you pay the associated transaction costs. In addition, only buyers and households who refinance have to comply with the LTV and PTI requirements, while other homeowners have to adhere to the minimum payment requirement of the amortization schedule. The solution to the household problem is given by

$$\begin{aligned}
V_j(z,x,h,m) = \max \Big\{ & V_j^R(z,x,h,m), V_j^B(z,x,h,m)\\
& V_j^S(z,x,h,m), V_j^{RF}(z,x,h,m) \Big\},
\end{aligned} \tag{3.10}$$

with the policy functions that maximize the Bellman equation for the chosen discrete tenure status

$$\Big\{ c_j(z,x,h,m), s_j(z,x,h,m), h'_j(z,x,h,m), m'_j(z,x,h,m), b'_j(z,x,h,m) \Big\}.$$

3.2.2 Rental market

There is a unit mass of homogeneous rental firms f that operate in a competitive market with free entry and exit. Rental firms offer rental housing to households, and are owned by foreign investors. The required rate of return of the investors is equal to the return on risk-free bonds r. The competitive rental rate p_r for a unit of rental housing is given by the user-cost formula,

$$p_r = \frac{1}{1+r}\Big[rp_h + \delta^r + \tau^h p_h\Big]. \tag{3.11}$$

Hence, the rental rate is such that it covers the cost of capital rp_h, the maintenance cost of the rental property δ^r, where $\delta^r > \delta^h$, and the property taxes $\tau^h p_h$.[8] Since the operating expenses are realized in the next period, these costs are discounted at the required rate of return of the investors.

3.2.3 Government

The main role of the government in the model is to tax households and rental firms, and provide social security benefits to those in retirement. Overall, the government runs a surplus, which it spends on activities that do not affect the other agents in the economy.

The government collects property taxes from the rental firms, and taxes the households using three different taxes, as described in equation (3.8). The labor income tax is modeled using a non-linear tax and transfer function $T(\tilde{y})$, as in Heathcote et al. (2017). This function is continuous and convex, and is meant to proxy for the progressive federal earnings taxes in the U.S.

$$T(\tilde{y}) = \tilde{y} - \lambda \tilde{y}^{1-\tau^p}, \tag{3.12}$$

where λ governs the level of the income tax, and τ^p controls the degree of progressivity. The argument $\tilde{y}$ is taxable labor income, which consists of labor income or social security benefits, net of deductions. If beneficial, a household deducts mortgage interest payments and property taxes before paying labor income taxes. Thus, we include some of the main features of the U.S. tax code with respect to housing; that is, imputed rents are not

[8] The assumption that rental property requires higher maintenance costs than owned housing is motivated by the potential moral hazard problem of rental housing. This is also a common feature of housing models to generate a benefit of owning compared to renting a house (see, e.g., Piazzesi and Schneider (2016)).

taxed, mortgage interest payments and property taxes are tax deductible, and labor income after deductions is subject to a progressive tax schedule.

3.3 Calibration

The model is calibrated to the U.S. economy. As our aim is to capture a steady state of the economy, we conduct the calibration using long-run averages of parameter values and moments. As this class of models has a hard time matching the strong skewness in wealth that we see in the data, we choose to focus on the bottom 90 percent of the population in terms of net worth. In this paper, we are interested in how households' consumption responses to shocks are affected by different policies in the mortgage and housing markets. Households with very high levels of wealth are likely to be unconstrained in their spending, and their responsiveness to shocks will presumably not depend much on frictions in mortgage and housing markets. Thus, restricting our attention to the bottom 90 percent of the wealth distribution should not materially affect our findings.

3.3.1 Independently calibrated parameters

Most of the parameters are calibrated independently, either computed from data or taken directly from other studies. These parameters are listed in Table 3.1. In the next section, we move on to estimate the remaining parameters using simulated method of moments.

Demographics and preferences

Households enter the model economy at age 23. At age 65, all households retire, and by age 83 all households have exited the economy. Before retirement, households do not face a risk of dying, but in between age 65 and 82 the probability of surviving to the next period ϕ_j is taken from the Life Tables for the U.S., social security area 1900-2100, for males born in 1950 (see Bell and Miller (2005)).

The coefficient of relative risk aversion σ in the utility function is set to 2, in line with much of the literature. The age-dependent utility shifter e_j, which captures how household size changes with the life cycle, is calibrated from the Panel Study of Income Dynamics (PSID), survey

Parameter	Description	Value
σ	Coefficient of relative risk aversion	2
τ^{ss}	Social security tax	0.153
τ^{h}	Property tax	0.01
r	Interest rate, bonds	0
r^{m}	Interest rate, mortgages	0.036
θ	Down-payment requirement	0.10
ψ	Payment-to-income requirement	0.28
δ^{h}	Depreciation, owner-occupied housing	0.03
ς^{I}	Home insurance	0.005
ς^{b}	Transaction cost if buying house	0.025
ς^{s}	Transaction cost if selling house	0.07
R	Replacement rate for retirees	0.5
B^{max}	Maximum benefit during retirement	60.4

Table 3.1: Independently calibrated parameters, taken from the data and other studies

Note: Where relevant, the parameter values are annual. The maximum benefit during retirement B^{max} is stated in 1000's of 2018 dollars.

years 1970 to 1992. Specifically, we estimate e_j with a regression of family size on a third-order polynomial of age, and then take the square root of the predicted values.

Taxes

Based on Harris (2005), the social security tax τ^{ss} is set to 15.3 percent of earnings, which corresponds to the total payroll tax for both employers and employees. The property tax rate τ^{h} is taken from the 2009, 2011, and 2013 waves of the American Housing Survey (AHS). The median real estate tax as a share of the housing value is approximately 1 percent.

Bonds, housing and mortgages

Using yearly data from 1997 to 2013 on 3-month Treasury bill rates, deflated by the Consumer Price Index (CPI), the mean real rate is 0.06 percent.[9] The interest rate on risk-free bonds is therefore set to zero. The average real interest rate on long-term mortgages for the same period is equal to 3.6 percent. This is computed from the Federal Reserve's series

[9]We use data from the Federal Reserve Bank of St Louis of the 3-month Treasury bill rate from the secondary market, seasonally adjusted, and the CPI data is the U.S. city average CPI for all urban consumers, all items.

of the contract rate on 30-year fixed-rate conventional home mortgage commitments, deflated by the CPI. Hence, we choose a yearly mortgage interest rate of 3.6 percent.

Between 1976 and 1992, the average down payment of first-time buyers in the U.S. ranged from 11 to 21 percent of the house value (U.S. Bureau of the Census, Statistical Abstract of the United States (GPO), 1987, 1988, and 1994). We use the lower bound of this interval, and set the down-payment requirement θ for new mortgages to 10 percent, as this helps us capture the upper tail of the LTV distribution. The payment-to-income requirement ψ is set to 0.28, consistent with Greenwald (2018). The depreciation rate of owned housing is taken from Harding et al. (2007) who estimate the median depreciation rate of owned housing, gross of maintenance, to be 3 percent. The transaction costs for buying and selling a house are set to 2.5 and 7 percent of the house value, respectively. These values are taken from Gruber and Martin (2003). The home insurance rate ς^I is set to 0.005 percent of the house value, which is roughly in line with the median property insurance payment in the 2013 AHS.

Initial assets

To match the distribution of wealth and the correlation between earnings and wealth among the young, we distribute initial assets a to the newborn cohort in the model similarly to Kaplan and Violante (2014). In the model, we divide newborns into 21 equally-sized groups based on their earnings. The probability of being born with initial assets and the amount of these assets vary across earnings bins. These probabilities and amounts are estimated based on data from the Survey of Consumer Finances (SCF). Specifically, we divide households of age 23-25 in the SCF for survey years 1989 to 2013 into 21 equally-sized groups based on their reported earnings. We assume that a household has positive initial assets in the data whenever its asset holdings are larger than 1,000 in 2013 dollars. Within each earnings bin, we then compute the share of households that meet this requirement and the median net worth of these households. For each bin, we scale the median net worth by median earnings for the working-age population in the data. We rescale by median earnings in the model when we allocate the initial assets to households in the model economy.

Labor income

The labor income process is inspired by Cocco et al. (2005). There is an age-dependent and a household-specific component of earnings. Further, households of working age face permanent and transitory earnings risk, while households in retirement only experience transitory shocks to their social security benefits. The estimation of the earnings process is described in detail in Appendix 3.C.

Log earnings for household i of age j are given by

$$\log(y_{ij}) = \alpha_i + g(j) + n_{ij} + \nu_i \quad \text{for } j \leq J_{ret}, \tag{3.13}$$

where α_i is the household fixed effect, distributed $N(0, \sigma_\alpha^2)$, and $g(j)$ is the age-dependent component of earnings, which captures the hump-shaped life-cycle profile. n_{ij} is an idiosyncratic random-walk component, which evolves according to a permanent income shock η_{ij}, distributed $N(0, \sigma_\eta^2)$. The household also draws an i.i.d. transitory shock ν_i, distributed $N(0, \sigma_\nu^2)$, which is uncorrelated with the permanent earnings shock. The log of the permanent earnings state z_{ij} in the model is given by the sum of the household-fixed component, the age-dependent component of earnings, and the random-walk component, i.e., $\log(z_{ij}) = \alpha_i + g(j) + n_{ij}$.

The social security benefits in retirement are given by a fixed proportion R of permanent earnings in the period before retirement, subject to a cap B^{max}. The common replacement rate R is taken from Díaz and Luengo-Prado (2008) and is set to 50 percent, whereas B^{max} is computed from Social Security Administration data. Further, the benefits are affected by transitory shocks, drawn from the same distribution as the transitory earnings shocks. Formally,

$$\log(y_{ij}) = \min\left(\log(R) + \log(z_{i,J_{ret}}),\ \log(B^{max})\right) + \nu_i \quad \text{for } j > J_{ret}. \tag{3.14}$$

To estimate equation (3.13), we use PSID data from the survey years 1970 to 1992. In the estimation of the age-dependent components of earnings $g(j)$, we follow Cocco et al. (2005). We estimate the variances of the permanent and transitory shocks as in Carroll and Samwick (1997). The variance of the fixed-effect shock is estimated as the residual variance in earnings of the youngest cohort, net the deterministic trend value and the variances of the permanent and the transitory shocks. The estimated variances of the earnings shocks are displayed in Table 3.2.

Parameter	Description	Value
σ_α^2	Fixed effect	0.156
σ_η^2	Permanent	0.012
σ_ν^2	Transitory	0.061

Table 3.2: Estimated variances of earnings shocks
Note: Household earnings contain a fixed household component. Throughout working life, earnings are subject to permanent and transitory shocks, while in retirement there is only transitory earnings risk. Estimated with PSID data, years 1970 to 1992.

3.3.2 Estimated parameters

The parameters that are estimated to match a set of data moments are listed in Table 3.3. Unless otherwise noted, we use data from the SCF, pooled across the 1989 to 2013 survey years. All parameters in Table 3.3 are jointly estimated, taking the independently calibrated parameters in Table 3.1 as given.[10]

Parameter	Description	Value	Target moment	Data	Model
α	Consumption weight in utility	0.80	Median house value-to-earnings, age 23–64	2.26	2.26
β	Discount factor	0.956	Mean net worth, over mean earnings age 23–64	1.38	1.38
υ	Strength of bequest motive	5.60	Net worth mean age 75 over mean age 50	1.64	1.64
δ^r	Depreciation rate, rentals	0.076	Homeownership rate, age 23–35	0.44	0.44
$\underline{\text{h}}$	Minimum owned house size	199	Homeownership rate, all ages	0.67	0.67
ς^r	Refinancing cost	2.77	Refinancing share, homeowners	0.08	0.08
λ	Level parameter, tax system	1.69	Average marginal tax rates	0.13	0.13
τ^p	Progressivity parameter	0.14	Distribution of marginal tax rates	N.A.	N.A.

Table 3.3: Estimated parameters
Note: Parameters estimated using simulated method of moments. The first two columns list the parameters and their descriptions. The third column shows the estimated parameter values. The fourth column contains the descriptions of the targeted moments, while column five lists their respective values in the data. Finally, the last column states the values of the corresponding model moments, achieved by using the parameter values in column three. The minimum owned house size $\underline{\text{h}}$ and the fixed refinancing cost ς^r are in 1000's of 2018 dollars.

The consumption weight in the utility function α controls the share of expenditures that is allocated to consumption versus housing services. This weight is set to 0.80 to match the median house value-to-earnings

[10]When we solve the baseline model, the housing supply is chosen such that the price of a unit of owned housing is equal to the price of a unit of consumption, i.e., $p_h = 1$. In turn, the rental rate is given by equation (3.11). See the Appendices for a detailed description of the solution method and the equilibrium definition.

ratio, among the working-age homeowners. The discount factor β affects the savings decisions. It is therefore used to match the mean net worth over mean earnings, among households of age 23 to 64. The resulting yearly discount factor is 0.956. To capture the strength of the bequest motive, the utility shifter of bequests υ is used to match the mean net worth of households aged 75 over the mean net worth of households aged 50. The parameter value is estimated to be 5.60.

The decision to buy a house instead of renting housing services is affected by a number of factors in the model. Abstracting from frictions in the mortgage and housing markets, households generally prefer to own. This positive net benefit of owning is partly due to the preferential tax treatment of owned housing, i.e., mortgage interest payments and property taxes are tax deductible and imputed rents are left untaxed. However, because there are frictions in the mortgage and housing markets, an additional benefit of owning is required to incentivize households to buy when they are young. Therefore, we estimate the depreciation rate of rental housing δ^r to match the homeownership rate among young households, aged 23 to 35. The depreciation rate needed to meet this target is 7.6 percent. The minimum house size available for purchase $\underline{\mathrm{h}}$, which is strictly larger than the minimum house size available for rent, is set to match the overall homeownership rate in the data. To capture the liquidity of housing equity, we estimate the fixed refinancing cost ς^r. With a cost slightly below $2,800$ in 2018 dollars, we match the 8 percent refinancing rate among homeowners as stated in Chen et al. (2020).

The two parameters of the tax and transfer function $T(\tilde{y})$ are estimated to match the level and the progressivity of earnings taxes in the U.S. The level parameter λ is set to 1.69, to match the average marginal earnings tax rate after deductions among the working-age population. The progressivity of the earnings tax is controlled by parameter τ^p. This parameter is set to 0.14, to minimize the sum of the absolute difference between the fraction of households exposed to the different statutory tax brackets in the data compared to the model. Since the tax schedule is continuous in the model, households are allocated to their nearest tax bracket in the data for this calibration exercise. The data on tax rates is taken from Harris (2005).

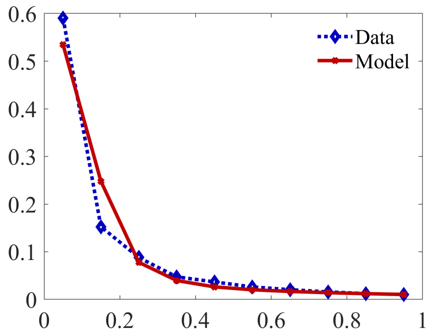

(a) Distribution, liquid savings-to-earnings

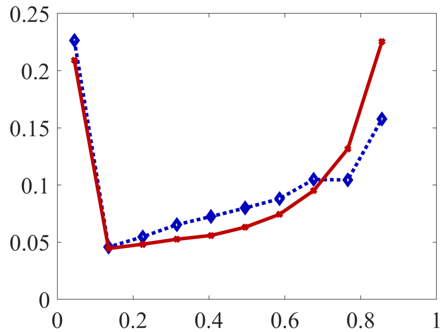

(b) Distribution, loan-to-value

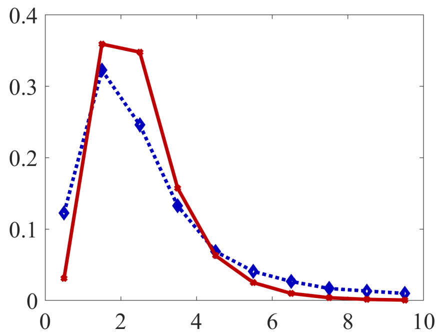

(c) Distribution, house value-to-earnings

Figure 3.1: Comparison of data versus model: non-targeted distributions
Note: The data is from the SCF, survey years 1989-2013. The model refers to the baseline economy. In Figure 3.1a and Figure 3.1c, only working-age households are included, and Figure 3.1b only displays homeowners.

3.3.3 Data versus model: distributions

At the heart of our research question is the need for the model to capture the extent to which households are constrained. Households may be constrained in their spending if they have low levels of liquid bond savings. How constrained a homeowner is also depends on how much equity is available in the house, and if increased mortgage financing is possible. In Figure 3.1, the distributions of liquid savings-to-earnings, LTVs, and house value-to earnings are shown for the model and for data from the

SCF.[11] Further, the life-cycle profiles of LTV and homeownership inform us about who the constrained homeowners are. Housing and mortgage choices are tightly linked to the age of households, as seen in Figure 3.2.

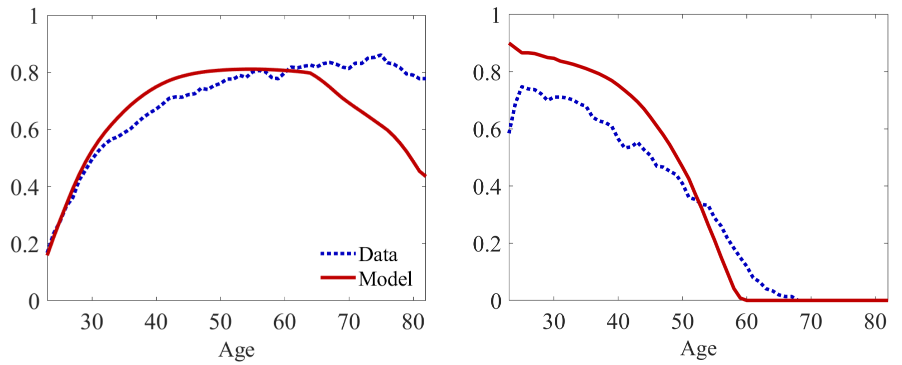

(a) Homeownership rate **(b)** Median LTV

Figure 3.2: Comparison of data versus model: non-targeted life-cycle profiles
Note: The data is from the SCF, survey years 1989-2013. The model refers to the baseline economy. The median LTV is computed among homeowners.

3.4 Results

Equipped with our model, we now turn to the quantitative analysis. We start by carefully analyzing the determinants of MPCs in our baseline model. Then, we consider how permanent and temporary changes in LTV and PTI requirements affect individual and aggregate consumption responses to wealth shocks. In the case of temporary policies, we complement the analysis by solving for optimal policies and investigate how they vary depending on the magnitude of the wealth shocks.

We define the marginal propensity to consume for household i of age j as

$$MPC_{ij} \equiv \frac{c_{ij}(z, x + \Delta_x, h, m) - c_{ij}(z, x, h, m)}{\Delta_x}, \tag{3.15}$$

[11]We define liquid savings in the SCF as the sum of cash, checking, savings, money market, and call accounts, prepaid cards, directly-held mutual funds, stocks, and bonds, less any credit card debt balance. Cash is assumed to be five percent of the balance in the variable *liq* in the SCF, similar to Kaplan and Violante (2014). We define net worth to be the sum of liquid savings and housing wealth less mortgages.

where $c_{ij}(z, x, h, m)$ is consumption for household i of age j if there is no shock, and $c_{ij}(z, x + \Delta_x, h, m)$ is consumption when there is a shock of size Δ_x. Intuitively, the MPC is the fraction of the shock Δ_x that is spent on non-housing consumption. The unexpected change in cash-on-hand Δ_x is referred to as a wealth shock. This shock is meant to capture a change in available resources that could stem from various sources, such as unexpected changes in asset prices or labor income.[12]

As more stringent lending standards are often introduced to alleviate the costs of large shocks in the economy, Δ_x will take on sizable values in our experiments. When subject to larger shocks, some households may want to change their discrete tenure choice. We refer to these households as *switchers*, whereas households who do not change their discrete choice are referred to as *non-switchers*. For example, a household is a switcher if it were to have been a renter, but chooses to become a homeowner due to the wealth shock.

3.4.1 Dissecting MPCs in a housing model

Before we study the impact of stricter mortgage lending standards, it is useful to understand the underlying determinants of MPCs in the model. We begin by showing the MPCs of a negative wealth shock of $1,000$ dollars.[13] Later, we also explore how the MPC varies with the sign and magnitude of the shock.

Figure 3.3 shows that there is considerable heterogeneity in MPCs across households. At the right-hand tail, there is a large group of households that have an MPC of one, and thus reduce their spending one-for-one with the fall in cash-on-hand. They are so-called hand-to-mouth households. In contrast, other households increase their non-housing consumption in response to the negative shock, which implies that their MPCs are negative. In between these extremes, there is a significant mass over the whole support.

---

[12] We think of a negative wealth shock as representing an economic downturn, though admittedly a stylized one.

[13] Hereafter, dollars refer to 2018 dollar value.

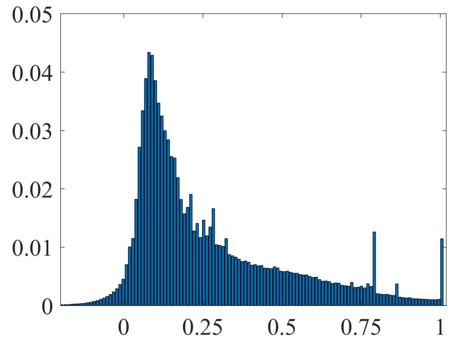

Figure 3.3: Distribution of MPCs
Note: Wealth shock of $-1,000$ dollars.

To gain further intuition about the distribution of MPCs, we first consider three groups of non-switchers, i.e., those who do not change their discrete choice in response to the shock. The first group consists of *renters*. We call the second group *constrained owners*, which we define as owners who choose an LTV above 0.8 and/or follow the mortgage repayment plan in the absence of the wealth shock. The last group, *unconstrained owners*, comprises households who choose an LTV below 0.25 and a mortgage level below that implied by the amortization plan in the case when there is no shock. Clearly, there are households that do not fall into either of these groups. The chosen groups are only meant to illustrate the key determinants of MPCs.

Figure 3.4a shows how MPCs depend on the ratio of liquid savings to earnings that households would choose if there was no shock. Naturally, households that expect to hold considerable amounts of liquid bonds are better prepared to handle negative shocks and thus have lower MPCs. For renters and constrained owners, lower bond holdings signal that these households were already constrained before the shock. When hit by a negative wealth shock, they respond by decreasing non-housing consumption. Renters with no savings (poor hand-to-mouth) rent in a frictionless rental market, so their drop in non-housing consumption equals the consumption expenditure share $\alpha \approx 0.8$. This explains the spike around 0.8 in Figure 3.3. Constrained owners, with low levels of liquid savings (wealthy hand-to-mouth), cannot freely access their housing equity. As a consequence, they respond by reducing non-housing consumption and have MPCs around one. These households thus comprise the right-hand tail in Figure 3.3. The MPCs of unconstrained owners remain relatively moderate even for low levels of liquid assets-to-earnings. These households expect to pay off more on their mortgage than what is stipulated by their amortization plan, and can thus adjust by paying off less in response to the shock.

In Figure 3.4b, we show that households with a higher transitory income tend to have lower MPCs. This observation complements the findings in Figure 3.4a. Households with a high transitory income component are more likely to save in order to smooth consumption over time. Thus, when hit by a negative wealth shock, these households have the possibility to save less than planned. Households with a low transitory shock are relatively poor today and expect higher earnings in the future. Therefore, they want to save little to begin with, and respond strongly to the negative wealth shock by consuming less. Again, the MPC of unconstrained owners is generally lower.

A key feature of Figure 3.3 that we have not discussed thus far is the large portion of households with an MPC of around 0.1. Our results in Figure 3.4a and Figure 3.4b indicate that these are households with a high transitory income and/or those who can use their liquidity buffer to cushion the negative wealth shock. Thus, these households are fairly unconstrained in their spending.

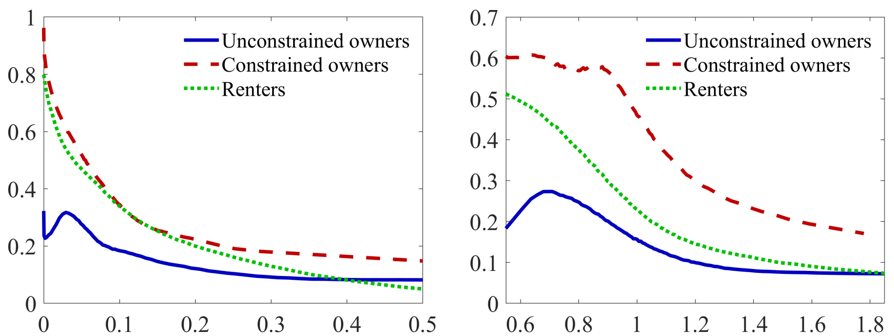

(a) MPC across liquid savings-to-earnings

(b) MPC across transitory earnings

Figure 3.4: Decomposing the mean MPC of non-switchers
Note: MPCs for working-age households from a wealth shock of $-1,000$ dollars

The households who change their discrete choice, i.e., the switchers, behave quite differently from the non-switchers described above. Almost all switchers have sizable negative MPCs, most of them much lower than what is shown in Figure 3.3. On average, their MPC is approximately -8. As the group of switchers account for less than one percent of the population in the case of a wealth shock of $-1,000$ dollars, the mean MPC in the economy is still relatively high and equal to 0.19.

For a negative wealth shock, there are two important groups of switchers. The first group consists of households who choose to abstain from buying a house due to the shock. These households are, on average, younger and have a lower income than other buyers. Although their total spending may decrease due to the wealth shock, their non-housing consumption increases as they avoid paying the down payment and the transaction cost of buying. Out of all households that would buy a house in the absence of the wealth shock, 4.1 percent of them decide not to.

The second group of switchers comprises households who choose to refinance their mortgage instead of following their amortization plan, due to the negative wealth shock. They have illiquid housing wealth that they access by paying the refinancing cost. As the refinancing cost is sizable, it only makes sense for households in dire need of liquidity to pay the cost. Households who choose to refinance, due to the shock, only make up one percent of all initial stayers, and they tend to have a low transitory income. Once these households access their housing equity, they significantly increase their consumption.

In Figure 3.5, we decompose the effects of non-switchers and switchers for the mean MPC across shock sizes. Figure 3.5a shows that the average MPC of non-switchers is close to 0.3 for most shocks, although the MPC is falling somewhat for larger positive shocks as households become increasingly unconstrained. Clearly, the MPC of switchers, as depicted in Figure 3.5b, differs remarkably from that of non-switchers. For smaller wealth shocks, the MPC is very low. As the shocks become more significant, the MPC becomes less negative. When households change a discrete choice, this leads to a jump in non-housing consumption. Contingent on switching, the absolute size of the jump in consumption largely depends on the level of the down-payment requirement and the transaction costs of buying and refinancing. For example, the savings from not paying the down payment and the transaction cost of buying do not depend on the shock size, for a household who abstains from buying. The lower the transaction costs are, the lower is the change in consumption.[14]

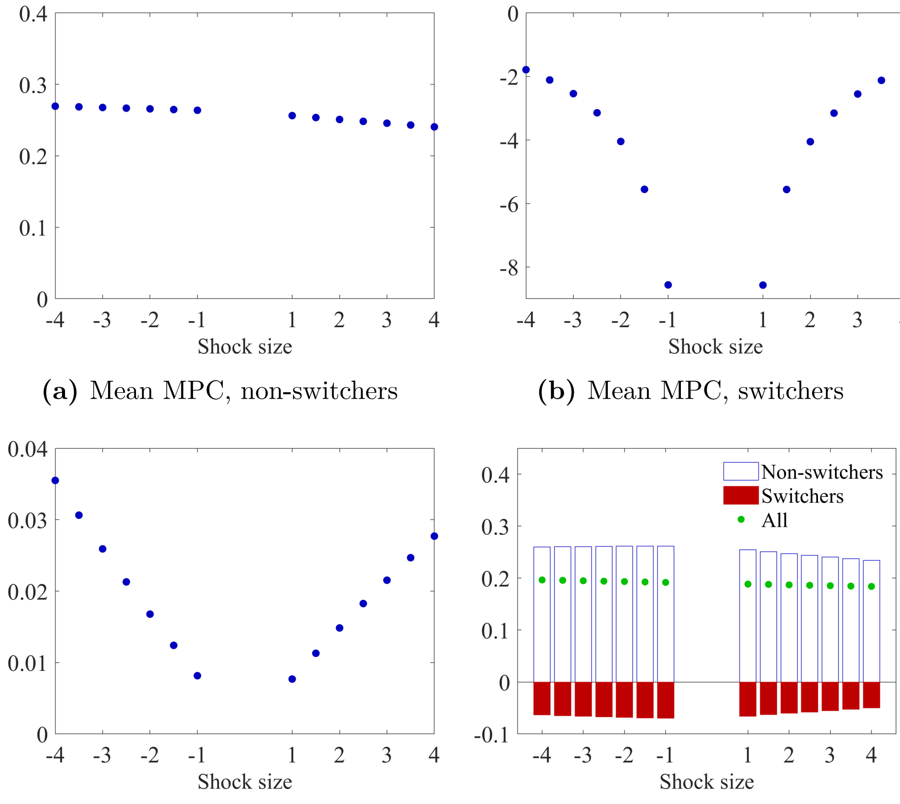

(a) Mean MPC, non-switchers
(b) Mean MPC, switchers
(c) Share who switches
(d) Mean MPC, all

Figure 3.5: Decomposing the mean MPC across shock size (thousands of dollars)
Note: Switchers are those who change their discrete choice in response to a shock.

Despite that the average MPC of switchers is sensitive to the shock size, Figure 3.5d shows that the mean MPC among all households is close to 0.19 for the range of shock sizes we consider. There are two reasons for this result. First, the fraction of switchers increase in the magnitude of the wealth shock, as seen in Figure 3.5c. Thus, even if the MPC of

[14]See Appendix 3.D.1 for a comparison of the average MPCs of switchers in a setting where there are no refinancing costs or no transaction costs for buying and selling a house.

switchers becomes less negative for larger shocks, the extensive margin acts as a counter weight. Second, the fraction of switchers grows faster for negative than for positive wealth shocks. This off-sets the slight fall in MPCs among non-switchers as the shock becomes larger and positive.

3.4.2 Permanent changes in LTV and PTI

As shown in the previous section, there is a significant heterogeneity in MPCs, which arises due to costs and constraints in the housing and mortgage markets. Constrained homeowners are among the households with particularly high consumption responses to wealth shocks. Their debt levels are considerable and they generally have a limited access to liquid funds. As such, policymakers may find it reasonable to introduce stricter lending requirements. After all, higher debt levels are associated with higher MPCs.

A natural argument against stricter requirements is that they strengthen the financial frictions in the economy. By making it more difficult to borrow, the ability to smooth consumption in response to a wealth shock may worsen, causing an increase in MPCs. Moreover, one has to take into consideration the behavioral responses by households. The distribution of asset holdings is bound to change in response to new regulatory requirements. For example, a household that chooses to hold less debt due to a stricter LTV requirement may also choose to hold less liquid bonds now that it has more housing equity. Ultimately, the question of how mortgage lending standards affect consumption dynamics requires a quantitative analysis.

In this section, we study how the aggregate consumption response to a wealth shock and the distribution of MPCs across households change as a result of tougher LTV and PTI regulations. To quantify the effects of stricter policies on MPCs, we consider two relatively large changes. In the first experiment, we consider a permanent tightening of the LTV limit from 0.9 to 0.7. In the second experiment, the PTI requirement is 0.18 instead of the baseline value of 0.28. In both experiments, we solve for a new steady state, where we allow house prices to change under the assumption that the aggregate housing stock is fixed.[15]

[15] The pair of policies was chosen such that the percentage change in house prices is roughly the same.

The policies we consider impact the model economy in several important ways. Table 3.4 shows steady-state prices and moments across policies. When stricter regulations are in place, it is more difficult for households to buy houses. As a result, the homeownership rate is lower. Unsurprisingly, the policies reduce the average loan-to-value ratios in the economy. The mean net worth over mean earnings remains relatively stable, although it increases somewhat in the case of stricter LTV. In general, the LTV policy leads to larger changes in steady-state moments as compared to the PTI policy, even if the price effects are similar.

	Baseline	Stricter LTV	Stricter PTI
Max LTV	0.90	0.70	0.90
Max PTI	0.28	0.28	0.18
House price	1	0.965	0.959
Rent	0.086	0.086	0.086
Homeownership rate	0.674	0.605	0.647
Median house-to-earnings ratio	2.259	2.164	2.134
Mean net worth age 75 over 50	1.637	1.401	1.633
Median loan-to-value ratio	0.339	0.147	0.250
Mean net worth, over mean earnings	1.381	1.477	1.379
Mean liquid savings-to-earnings	0.752	0.765	0.765

Table 3.4: Steady-state prices and moments under permanent changes in lending policies

Although the debt levels are substantially reduced, the aggregate consumption response to wealth shocks and the distribution of MPCs are largely unaffected by the permanently stricter LTV and PTI policies. Figure 3.6a shows the aggregate consumption dynamics up to 10 years after a wealth shock of $-4,000$ dollars.[16] There are virtually no differences in the dynamics across policies. In Appendix 3.D.2, we show that this result holds for shock sizes of varying magnitudes and is independent of the sign of the shock. Moreover, Figure 3.6b shows that the distributions of MPCs are almost identical under all policies. These results are also robust to considerably larger changes in policies. A permanent change in the LTV limit to 0.5 or the PTI constraint to 0.1 produces very similar MPCs to those of the baseline model, as seen in Appendix 3.D.2. As there are no large changes in the distributions, there are also no significant changes in the role of switchers and non-switchers in the case of permanently stricter

[16]We assume that the shock is unexpected. To focus on the direct demand effect, we assume that prices are constant during the transition.

lending standards, see Appendix 3.D.2.

Overall, the behavioral responses of households are crucial for understanding why permanently stricter lending standards have such a small impact on MPCs. When considering permanent policies in steady state, households are free to re-optimize, taking into account the new regulatory environment. How much households save in liquid assets is driven by their desire to insure against negative earnings shocks. The amount of precautionary savings is governed by deep parameters, e.g. the risk-aversion parameter σ, rather than lending standards set by the government. As such, there are only small differences in liquid bond holdings across policies, as indicated by the mean liquid savings-to-earnings ratio in Table 3.4.

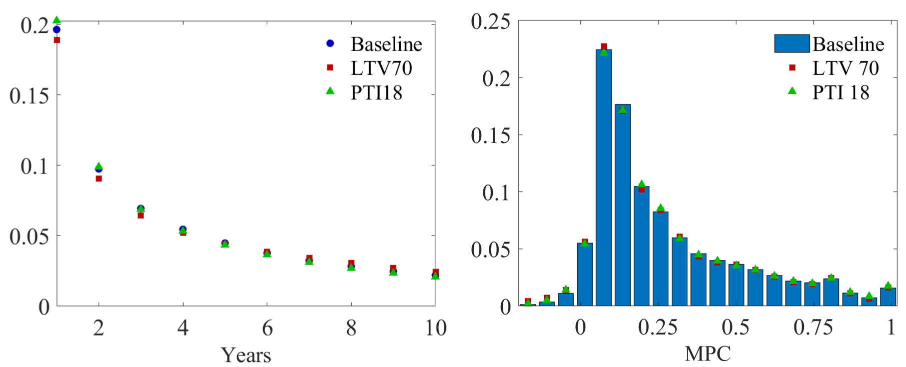

(a) Mean MPC over time **(b)** Distribution of MPCs in $t = 1$

Figure 3.6: MPCs for different permanent policies
Note: MPCs from a wealth shock of $-4,000$ dollars in $t = 1$. In the baseline model, the LTV limit is 90 percent and the PTI constraint is 28 percent.

3.4.3 Temporary changes in LTV and PTI

Can temporary changes in LTV and PTI affect consumption responses?

A key conclusion from the previous section is that permanent policies appear to have a limited ability to affect consumption responses to wealth shocks. We now move on to analyze whether temporary policies can more effectively impact households' MPCs. Just like in the analysis of permanent policies, we begin by studying a wealth shock of $-4,000$ dollars. We let this wealth shock occur in time period $t = 2$. The shock is not expected by households, but we make the strong assumption that a hypothetical regulatory authority has perfect foresight. In an attempt to cushion the negative consumption response in $t = 2$, a stricter credit policy is enforced in $t = 1$, but then returns to its baseline value in $t = 2$. The policy is unexpectedly implemented in $t = 1$, but households know with certainty that the lending standards will be back to normal in the next period.

The main role of the temporary policy is to reallocate consumption over time. Since we abstract from price changes in this part of the analysis, cumulative consumption over time will be largely independent of whether there is a policy in place or not. Thus, the temporary policy may only be effective at dampening the consumption response in $t = 2$ if it can lower the spending in $t = 1$.

Qualitatively, the aggregate consumption effect in $t = 1$ is ambiguous. The policy affects households who would otherwise choose larger mortgages than what is allowed under the new policy. Thus, only households who refinance or buy a house in the absence of the policy are potentially affected. The group of households who would refinance without the policy lower their consumption in response to the policy for two reasons. First, households who refrain from refinancing cut back on consumption as they no longer extract any housing equity. Second, those who continue to refinance also need to reduce their consumption as the amount of equity extraction is restricted by the policy. Furthermore, households who continue to buy a house need to finance their home with more equity and thus decrease consumption. Households who abstain from buying a house, however, increase their consumption since they no longer have to finance the down payment or pay any transaction costs.

Quantitatively, the consumption responses in $t = 2$ are dampened as a result of the temporary stricter lending standards. Figure 3.7a compares the consumption dynamics of the baseline model where there is no policy change to the case where the LTV limit is lowered to 0.7 in $t = 1$. Contrary to the results for permanent policies, the aggregate MPC out of the negative wealth shock is considerably reduced on impact ($t = 2$), and stays below the no-policy case for several years. The muted consumption response is made possible as the temporary stricter LTV requirement makes households cut consumption in $t = 1$.[17]

Figure 3.7b shows that a temporary change in the PTI limit can also reduce the consumption response in $t = 2$, although this policy appears somewhat less effective at achieving this goal. It is important to note, however, that it is also possible to get strong consumption responses from a temporary change in PTI. In results that we do not report, a temporary change in the PTI requirement to 0.1 leads to consumption responses that are quantitatively similar to reducing the LTV limit to 0.7.

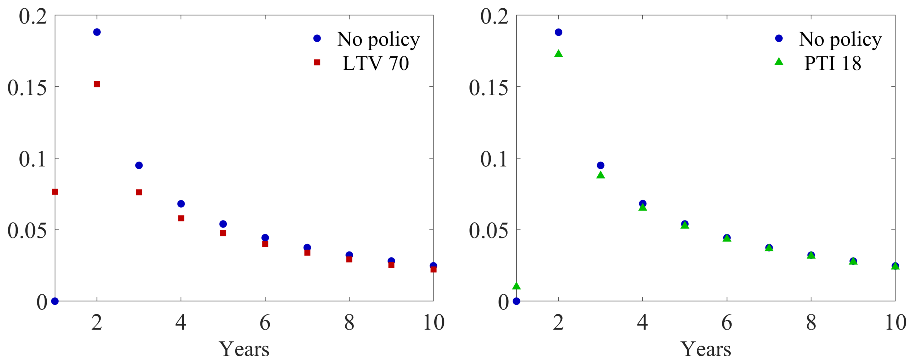

(a) Mean MPC over time **(b)** Mean MPC over time

Figure 3.7: MPCs for different temporary LTV and PTI policies
Note: Consumption responses under a temporary stricter policy in $t = 1$, that is reversed in $t = 2$. Unexpected wealth shock of $-4,000$ dollars in $t = 2$. The consumption responses are normalized by $-4,000$ dollars also in $t = 1$, where there is only a change in policy and no shock has occurred. In the baseline model with no temporary policy, the LTV limit is 90 percent and the PTI constraint is 28 percent.

Can temporary policies be welfare improving?

Although temporary policies may successfully dampen the consumption response to a negative wealth shock, it is not obvious whether and under what circumstances temporary policies improve welfare. On the one hand, households may benefit from the policy as it makes them increase their savings, thus making them better prepared to face the wealth shock. On the other hand, any fall in consumption in $t = 1$ reduces welfare in that period. Also, households may already save sufficiently for precautionary reasons. If the policy makes households save more than necessary, it has

[17]Note that the fall in consumption in $t = 1$ shows up as a positive MPC in the figure, as the consumption response is normalized by the negative wealth shock.

a negative impact on welfare.

To better understand the welfare implications of temporary lending policies, we solve for optimal LTV and PTI requirements in $t = 1$. We define an optimal policy as a policy that maximizes the mean ex-post consumption equivalent variation (henceforth CEV). More specifically, for each household alive at $t = 1$, we compute the per-period percentage change in consumption under the no-policy scenario needed to make the household indifferent between a policy and no policy. Our welfare measure is then the mean of these household-specific CEVs.[18] We do not consider policies that are more lenient than the benchmark lending requirements.

We find that temporary policies can be optimal, but only if the bust is sufficiently large. For example, the optimal policy for the wealth shock of $-4,000$ dollars is to keep lending standards at baseline levels throughout. However, when we consider a more extreme case, where all households are exposed to a wealth shock of $-12,000$ dollars, a temporary stricter LTV limit of 0.86 is optimal.[19] Yet, it continues to be optimal to leave the PTI requirement untouched at 0.28.

At the optimal LTV level, the mean MPC in the bust period is only slightly reduced and the average welfare gain is small. The nearly negligible changes in aggregate consumption dynamics are shown in Appendix 3.D.3. In terms of welfare, we find that the mean CEV is 0.0004 percent under the optimal LTV policy.

One reason for the small average welfare effect is that a vast majority of households are unaffected by the policy, and thus have a CEV of zero. The welfare effects can be substantial at the household level. Figure 3.8 shows the mean CEV across labor income shocks in $t = 2$, for a temporary LTV policy of 0.86. We limit the sample to only include households that change their mortgage decision in response to the policy change. The filled markers correspond to the welfare effects of introducing the policy when an unexpected shock of $-12,000$ dollars follows, whereas the hollow markers indicate the welfare effects of implementing the policy when there is no shock.

[18] A more thorough description of the welfare measure is provided in Karlman et al. (2020).

[19] As this shock is very large, we assume that no household can end up with a cash-on-hand lower than the lowest grid point used in the baseline calibration. This corresponds to about $1,800$ dollars.

When there is a large bust, the policy is positive for households whose income realization is low. Intuitively, a household with an unlucky income draw in $t = 2$ benefits from the increased savings in $t = 1$. Figure 3.8a shows that households whose permanent income is about 20 percent lower than expected have a mean welfare gain of 0.2 percent. Similarly, Figure 3.8b shows that households with a very low transitory income have a mean welfare gain of more than 0.6 percent. As indicated by the hollow markers in Figure 3.8, the policy is mostly negative for households if there is no bust in $t = 2$.

The welfare costs of a temporary stricter policy can be considerable for households who experience better income draws in $t = 2$, even when the bust is large. These households are simply better equipped to handle the negative wealth shock. Thus, the costs of lower consumption in $t = 1$ outweigh any potential benefit from increased savings.

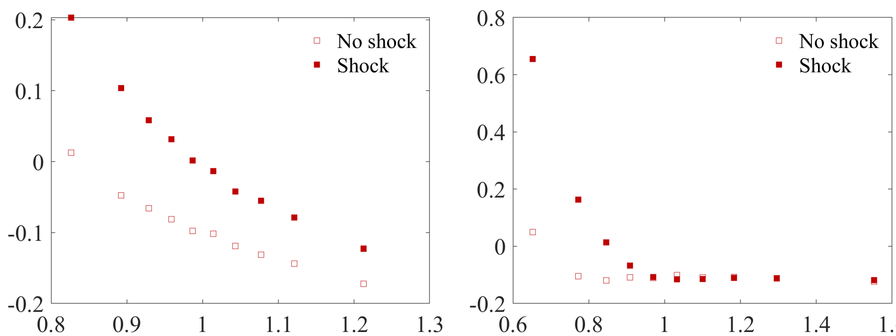

(a) CEV across permanent income **(b)** CEV across transitory income

Figure 3.8: Mean CEV (%) with or without wealth shock in $t = 2$
Note: The figures show the welfare effects of households that are directly affected by a temporary LTV policy of 86 percent in $t = 1$. The markers illustrate the mean welfare effect of ten equally sized groups, ordered by the variable on the x-axis. "No shock" refers to the welfare effects of introducing the policy when no subsequent wealth shock occurs. "Shock" refers to the welfare effects when a wealth shock of $-12,000$ dollars occurs in $t = 2$.

To shed some further light on the welfare effects, let us once more divide households into groups based on how they respond to the policy change. Recall that the policies only bind for households whose mortgage choice becomes limited by the new policy, i.e., refinancers and house buyers who would choose a larger mortgage absent the policy. Refinancers in $t = 1$ have usually drawn a very low transitory shock and are therefore in need of liquidity already in the first period. As a temporary stricter policy limits the extraction of housing equity in a period where liquid funds are valuable, these households have negative welfare effects on average. Households who continue to buy even after the policy has been introduced are also negatively affected on average. As more equity is needed to buy a house, their consumption drops in $t = 1$. Moreover, when they are hit by the negative wealth shock, they have a large fraction of their wealth in the illiquid housing asset and therefore find it difficult to smooth consumption. The only group that benefits from a temporary stricter policy are households who abstain from becoming homeowners in the boom. They increase their consumption in $t = 1$ and avoid being liquidity constrained in $t = 2$.

What are the effects of alternative shock scenarios?

There are alternative wealth-shock scenarios that are worth exploring. In particular, it can be argued that stricter LTV or PTI policies can be usefully implemented during a boom phase, as an exuberant economy may signal future busts.

To study the effects of including a boom period, we add a positive wealth shock of size Δ_x in $t = 1$, followed by a bust of the same magnitude in $t = 2$.[20] Figure 3.13 in Appendix 3.D.3 shows that temporary stricter LTV and PTI requirements continue to dampen the consumption responses in $t = 1$ and $t = 2$. Yet, for a given strictness of a temporary policy the consumption effect is lower, as the boom phase makes the policy less binding.

We find that the optimal policies are stricter when we consider a pronounced boom-bust episode, as compared to a scenario without a boom phase. For example, when the wealth shocks are of the size $12,000$ dollars, the optimal LTV and PTI policies are 0.8 and 0.18, respectively. Recall that with no boom phase, the optimal limits are 0.86 and 0.28. Why is that? First, during a boom there are fewer households who want to refinance and therefore the number of households who suffer from a stricter policy is lower. In the model, households who refinance often have a low transitory income. As the positive wealth shock in $t = 1$ is similar

[20] Admittedly, this example is highly stylized, but it still offers valuable insights into the effects of temporary policies.

to receiving a higher transitory income shock, fewer households find it optimal to tap into their housing equity. Second, when the bust is larger, the benefits from making households abstain from buying are greater.

When the boom-bust episode is more muted, the optimal policy is to leave the lending standards unchanged. For example, this is the case if we consider a boom of $4,000$ dollars followed by a bust of $-4,000$ dollars. When the boom is less strong, many households still want to refinance and thus the costs of stricter policies are larger. Furthermore, the benefit of keeping households from buying is reduced as the bust is less severe.

In the above analysis, we assume that the regulatory authority has perfect foresight and knows that there is a bust in $t = 2$. This informational advantage creates a rationale for the government to intervene. Clearly, this assumption is very strong. At the very least, we would expect there to be some noise in the government's signal about the future. Therefore, we also consider a case where there is a boom, but that no bust follows. Under this scenario, the optimal policy is to avoid temporary stricter policies. There is little to gain by restricting households from buying if there is no bust. Further, we consider a scenario where not only the policymaker but also the households have information about the coming bust. Also in this case, the optimal policy is to keep mortgage lending standards constant at current levels.

3.5 Concluding remarks

Since the Great Recession, policymakers in many countries have considered and implemented stricter mortgage lending standards. These policies aim at lowering household debt and, ultimately, reducing households' vulnerability to shocks. In this paper, we investigate if households' consumption responses to shocks depend on mortgage lending standards. Specifically, we study two types of policies in the mortgage market: stricter LTV and PTI requirements.

We find that permanently lower LTV and PTI limits reduce the debt level in the economy, but they are unsuccessful in dampening the aggregate consumption response to wealth shocks. In fact, the distribution of MPCs is only marginally affected by the permanently stricter policies. As the underlying incentives to insure against shocks are unchanged, households adjust their asset portfolio such that the more stringent borrowing requirements have little impact on their consumption sensitivity to shocks.

In contrast, we do find that temporary policies can dampen the consumption responses to shocks, but it does not come without any costs. Specifically, we find that LTV and PTI requirements introduced in a period before a downturn reduce the consumption fall during the bust. However, for such policies to be beneficial for households on average, strong assumptions about an informational advantage of the policymaker are needed, and the bust needs to be large.

There are a number of extensions to the analysis that would be worthwhile exploring in future work. First, in our analysis we abstract from propagation mechanisms through changes in prices or output, and focus on the immediate demand response from a wealth shock. A fruitful way forward would be to incorporate additional feedback effects of changes in demand to our framework. Arguably, households' direct endogenous responses to stricter mortgage regulations will be central even in a richer setting. Second, it would be interesting to see whether the results are generalizable to other types of shocks, such as changes to house prices.

References

Aastveit, K. A., E. Getz Wold, and R. Juelsrud (2020). Mortgage regulation and financial vulnerability. *Presented at the Nordic Junior Macro Seminar, May 27th.*

Angelini, P., S. Nicoletti-Altimari, and I. Visco (2012, November). Macroprudential, microprudential and monetary policies: conflicts, complementarities and trade-offs. Occasional Paper No. 140, SSRN.

Bell, F. C. and M. L. Miller (2005, August). Life tables for the United States Social Security Area 1900-2100. Actuarial study No. 120, Social Security Administration.

Boar, C., D. Gorea, and V. Midrigan (2020). Liquidity constraints in the U.S. housing market. *Working Paper*.

Carroll, C. D. and A. A. Samwick (1997). The nature of precautionary wealth. *Journal of Monetary Economics 40*(1), 41–71.

Chen, H., M. Michaux, and N. Roussanov (2020). Houses as atms: mortgage refinancing and macroeconomic uncertainty. *The Journal of Finance 75*(1), 323–375.

Cho, S.-W. and J. Francis (2011). Tax treatment of owner occupied housing and wealth inequality. *Journal of Macroeconomics 33*, 42–60.

Cocco, J. F., F. J. Gomes, and P. J. Maenhout (2005). Consumption and portfolio choice over the life cycle. *Review of Financial Studies 18*(2), 491–533.

Díaz, A. and M. J. Luengo-Prado (2008). On the user cost and homeownership. *Review of Economic Dynamics 11*(3), 584–613.

Druedahl, J. and T. H. Jørgensen (2017). A general endogenous grid method for multi-dimensional models with non-convexities and constraints. *Journal of Economic Dynamics and Control 74*, 87–107.

Ferrero, A., R. Harrison, and B. Nelson (2018, December). House price dynamics, optimal ltv limits and the liquidity trap. Discussion Paper No. DP13400, CEPR.

Floetotto, M., M. Kirker, and J. Stroebel (2016). Government intervention in the housing market: who wins, who loses? *Journal of Monetary Economics 80*, 106–123.

Gervais, M. (2002). Housing taxation and capital accumulation. *Journal of Monetary Economics 49*(7), 1461–1489.

Greenwald, D. (2018). The mortgage credit channel of macroeconomic transmission. *Working Paper*.

Gruber, J. W. and R. F. Martin (2003). Precautionary savings and the wealth distribution with illiquid durables. *International Finance Discussion Papers 773*.

Harding, J. P., S. S. Rosenthal, and C. Sirmans (2007). Depreciation of housing capital, maintenance, and house price inflation: estimates from a repeat sales model. *Journal of Urban Economics 61*(2), 193–217.

Harris, E. (2005, November). Effective marginal tax rates on labor income. A CBO paper, Congressional Budget Office.

Heathcote, J., K. Storesletten, and G. L. Violante (2017). Optimal tax progressivity: an analytical framework. *Quarterly Journal of Economics 132*(4), 1693–1754.

Kaplan, G., K. Mitman, and G. Violante (2020). The housing boom and bust: model meets evidence. *Journal of Political Economy*, Forthcoming.

Kaplan, G. and G. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica 82*(4), 1199–1239.

Karlman, M., K. Kinnerud, and K. Kragh-Sørensen (2020). Costly reversals of bad policies: the case of the mortgage interest deduction. *Working Paper*.

Lim, C. H., A. Costa, F. Columba, P. Kongsamut, A. Otani, M. Saiyid, T. Wezel, and X. Wu (2011, October). Macroprudential policy: what instruments and how to use them? Lessons from country experiences. IMF Working Paper No. 11/238.

Mian, A. and A. Sufi (2014). *House of debt: how they (and you) caused the Great Recession and how we can prevent it from happening again.* University of Chicago Press, Chicago.

Piazzesi, M. and M. Schneider (2016). Housing and macroeconomics. *Handbook of Macroeconomics 2*, 1547–1640.

Sommer, K. and P. Sullivan (2018). Implications of US tax policy for house prices, rents, and homeownership. *American Economic Review 108*(2), 241–74.

3.A Definitions of stationary equilibrium

Households are heterogeneous with respect to age $j \in \mathcal{J} \equiv \{1, 2, ..., J\}$, permanent earnings $z \in \mathcal{Z} \equiv \mathbb{R}_{++}$, mortgage $m \in \mathcal{M} \equiv \mathbb{R}_{+}$, owner-occupied housing $h \in \mathcal{H} \equiv \{0, \underline{\mathrm{h}}, ..., \bar{h} = \bar{s}\}$, and cash-on-hand $x \in \mathcal{X} \equiv \mathbb{R}_{++}$. Let $\mathcal{U} \equiv \mathcal{Z} \times \mathcal{M} \times \mathcal{H} \times \mathcal{X}$ be the non-deterministic state space with $\mathbf{u} \equiv (z, m, h, x)$ denoting the vector of individual states. Let $\mathbf{B}(\mathbb{R}_{++})$ and $\mathbf{B}(\mathbb{R}_{+})$ be the Borel σ-algebras on $\mathbb{R}_{++}$ and $\mathbb{R}_{+}$, respectively, and $P(\mathcal{H})$ the power set of $\mathcal{H}$, and define $\mathscr{B}(\mathcal{U}) \equiv \mathbf{B}(\mathbb{R}_{++}) \times \mathbf{B}(\mathbb{R}_{+}) \times P(\mathcal{H}) \times \mathbf{B}(\mathbb{R}_{++})$. Further, let $\mathbb{M}$ be the set of all finite measures over the measurable space $(\mathcal{U}, \mathscr{B}(\mathcal{U}))$. Then $\Phi_j(U) \in \mathbb{M}$ is a probability measure defined on subsets $U \in \mathscr{B}(\mathcal{U})$ that describes the distribution of individual states across agents with age $j \in \mathcal{J}$. Finally, denote the time-invariant fraction of the population of age $j \in \mathcal{J}$ by Π_j.

Stationary equilibrium, the baseline economy

Definition 1. A stationary recursive competitive equilibrium is a collection of value functions $V_j(\mathbf{u})$ with associated policy functions $\{c_j(\mathbf{u}), s_j(\mathbf{u}), h'_j(\mathbf{u}), m'_j(\mathbf{u}),$
$b'_j(\mathbf{u})\}$ for all j; prices $(p_h = 1, p_r)$; a quantity of total housing stock $\bar{H}$; and a distribution of agents' states Φ_j for all j such that:

1. Given the prices $(p_h = 1, p_r)$, $V_j(\mathbf{u})$ solves the Bellman equation (3.10) with the corresponding set of policy functions $\{c_j(\mathbf{u}), s_j(\mathbf{u}), h'_j(\mathbf{u}), m'_j(\mathbf{u}), b'_j(\mathbf{u})\}$ for all j.

2. Given $p_h = 1$, the rental price per unit of housing service p_r is given by equation (3.11).

3. The quantity of the total housing stock is given by the total demand for housing services[21]

$$\bar{H} = \sum_{\mathcal{J}} \Pi_j \int_U s_j(\mathbf{u}) d\Phi_j(U).$$

[21]We assume a perfectly elastic supply of both owner-occupied housing and rental units in the baseline steady state. This implies that supply always equals demand and thus we have market clearing.

4. The distribution of states Φ_j is given by the following law of motion for all $j < J$

$$\Phi_{j+1}(\mathcal{U}) = \int_U Q_j(\mathbf{u}, \mathcal{U}) d\Phi_j(U),$$

where $Q_j : \mathcal{U} \times \mathscr{B}(\mathcal{U}) \to [0,1]$ is a transition function that defines the probability that a household at age j transits from its current state $\mathbf{u}$ to the set $\mathcal{U}$ at age $j+1$.

Stationary equilibrium, after a permanent policy change

Definition 2. A stationary recursive competitive equilibrium after a permanent policy change is a collection of value functions $V_j(\mathbf{u})$ with associated policy functions $\{c_j(\mathbf{u}), s_j(\mathbf{u}), h'_j(\mathbf{u}), m'_j(\mathbf{u}), b'_j(\mathbf{u})\}$ for all j; prices (p_h, p_r); a quantity of total housing stock H; and a distribution of agents' states Φ_j for all j such that:

1. Given prices (p_h, p_r), $V_j(\mathbf{u})$ solves the Bellman equation (3.10) with the corresponding set of policy functions $\{c_j(\mathbf{u}), s_j(\mathbf{u}), h'_j(\mathbf{u}), m'_j(\mathbf{u}), b'_j(\mathbf{u})\}$ for all j.

2. Given p_h, the rental price per unit of housing service p_r is given by equation (3.11).

3. The housing market clears:

$$\begin{aligned} H &= \bar{H} \\ \text{where} \quad H &= \sum_{\mathcal{J}} \Pi_j \int_U s_j(\mathbf{u}) d\Phi_j(U) \end{aligned}$$

and $\bar{H}$ is the housing stock from the equilibrium of the baseline economy.

4. Distributions of states Φ_j are given by the following law of motion for all $j < J$

$$\Phi_{j+1}(\mathcal{U}) = \int_U Q_j(\mathbf{u}, \mathcal{U}) d\Phi_j(U),$$

3.B Computational method and solution algorithm

The computational method and the solution method are similar to those in Karlman et al. (2020). To summarize, we use the general generalization of the endogenous grid method G^2EGM by Druedahl and Jørgensen (2017) to solve for the value and policy functions. The number of grid points for permanent earnings N_Z, cash-on-hand N_X, housing sizes N_H, bonds-over-earnings N_B, and loan-to-value N_{LTV}, are 9, 140, 30, 25, and 41, respectively. The grid points are denser at lower levels of cash-on-hand and bonds-over-earnings. Further, we simulate 300 000 households for $J = 60$ periods.

3.C Labor income process

3.C.1 Data sample

Equation (3.13) is estimated using PSID data, survey years 1970 to 1992. The variable definitions and sample restrictions are the same as in Karlman et al. (2020).

3.C.2 Estimation

In this section, we describe how the exogenous earnings process in equation (3.13) is estimated. First, we estimate the deterministic life-cycle earnings profile $g(j)$, and then we move on to the variances of the fixed-effect component σ_α^2, the permanent shock σ_η^2, and the transitory shock σ_ν^2.

To estimate the deterministic age-dependent earnings component $g(j)$, we use yearly observations in the data for ages 20 to 64. Log household earnings $\log(y_i)$ are regressed on dummies for age (not including the youngest age), marital status, family composition (number of family members besides head and, potentially, wife), and a dummy for whether the household head has a college education. Household fixed effects are controlled for by running a linear fixed-effect regression. Finally, a third-order polynomial is fitted to the predicted values of this regression, which provides us with the estimate of the deterministic life-cycle earnings profile

$\hat{g}(j)$.

We follow Carroll and Samwick (1997) when we estimate the variances of the transitory (σ_ν^2) and permanent (σ_η^2) shocks. Define $\log(y_{ij}^*)$ as the logarithm of household i's earnings less the household fixed component $\hat{\alpha}_i$ and the deterministic life-cycle component.

$$\begin{aligned}\log(y_{ij}^*) &\equiv \log(y_{ij}) - \hat{\alpha}_i - \hat{g}(j)\\ &= n_{ij} + \nu_{ij} \qquad\qquad \text{for } j \in [1, J_{ret}],\end{aligned}$$

where the equality follows from equation (3.13). Define r_{id} as household i's d-period difference in $\log(y_{ij}^*)$,

$$\begin{aligned}r_{id} &\equiv \log(y_{i,j+d}^*) - \log(y_{ij}^*)\\ &= n_{i,j+d} + \nu_{i,j+d} - n_{ij} - \nu_{i,j}\\ &= n_{i,j+1} + n_{i,j+2} + ... + n_{i,j+d} + \nu_{i,j+d} - \nu_{i,j}.\end{aligned}$$

Since the transitory and permanent shocks are i.i.d., it follows that

$$\begin{aligned}\text{Var}(r_{id}) = \text{Var}(n_{i,j+1}) &+ \text{Var}(n_{i,j+2}) + ... + \text{Var}(n_{i,j+d})\\ &+ \text{Var}(\nu_{i,j+d}) + \text{Var}(\nu_{i,j})\\ = 2\,\sigma_\nu^2 &+ d\,\sigma_\eta^2.\end{aligned}$$

These variances are estimated by running an OLS regression of $\text{Var}(r_{id}) = r_{id}^2$ on d, including a constant term. The estimate of the variance of the permanent shock is given by the coefficient of d, and the estimate of the variance of the transitory shock is equal to the constant term divided by two. The estimate of the variance of the household fixed-effect component of earnings $\hat{\sigma}_\alpha^2$ is given by the residual variance in period $j = 1$,

$$\hat{\sigma}_\alpha^2 = \text{Var}\left(\log(y_{i1}) - \hat{g}(1)\right) - \hat{\sigma}_\eta^2 - \hat{\sigma}_\nu^2.$$

3.D Additional results

3.D.1 Baseline model

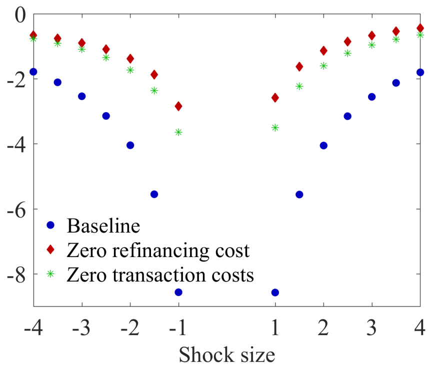

Figure 3.9: MPCs of switchers: no refinancing costs or transaction costs

Note: Mean MPC across shock size (thousands of dollars) among switchers, comparing the baseline model to a setting where there are no refinancing costs or no transaction costs for buying and selling a house. Switchers are those who change their discrete choice in response to a shock. For each new setting we solve for a new steady state, where we allow house prices to change under the assumption that the aggregate housing stock is fixed.

3.D.2 Permanent policies

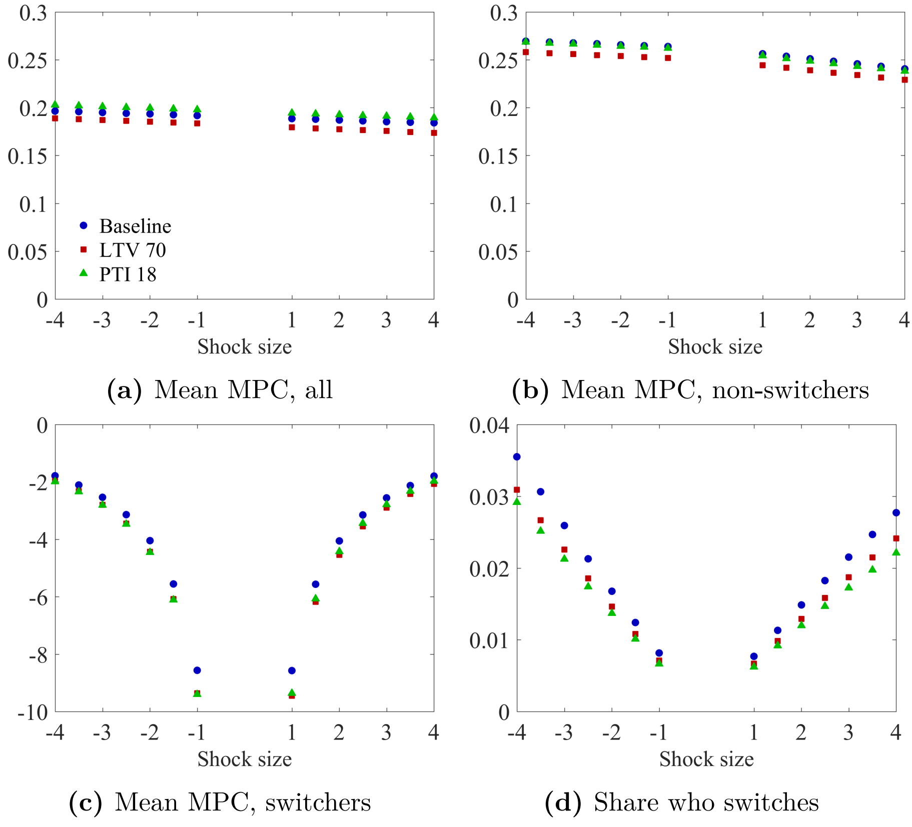

Figure 3.10: Decomposing the mean MPC across shock size (thousands of dollars)

Note: Switchers are those who change their discrete choice in response to a shock.

	Baseline	Stricter LTV	Stricter PTI
Max LTV	0.90	0.50	0.90
Max PTI	0.28	0.28	0.10
House price	1	0.893	0.846
Rent	0.086	0.085	0.085
Homeownership rate	0.674	0.527	0.568
Median house-to-earnings ratio	2.259	2.022	1.803
Mean net worth age 75 over 50	1.637	1.343	1.617
Median loan-to-value ratio	0.339	0.015	0.013
Mean net worth, over mean earnings	1.381	1.458	1.367
Mean liquid savings-to-earnings	0.752	0.790	0.803

Table 3.5: Steady-state prices and moments under permanent changes in the lending policies

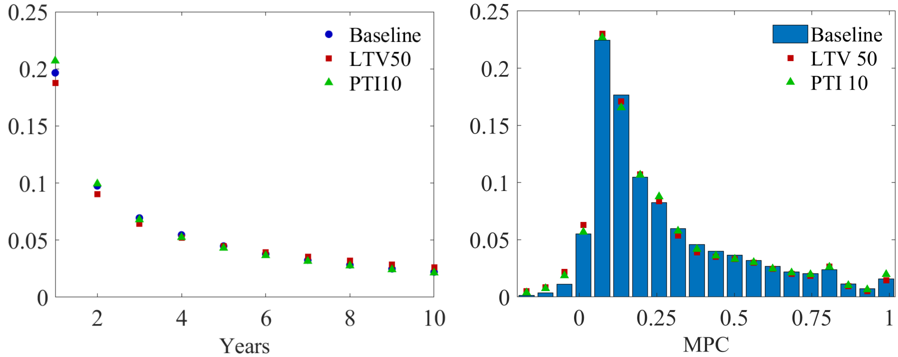

(a) Mean MPC over time **(b)** Distribution of MPC in $t = 1$

Figure 3.11: MPCs for different permanent policies
Note: MPCs from a wealth shock of $-4,000$ dollars in $t = 1$. In the baseline model, the LTV limit is 90 percent and the PTI constraint is 28 percent.

3.D.3 Temporary policies

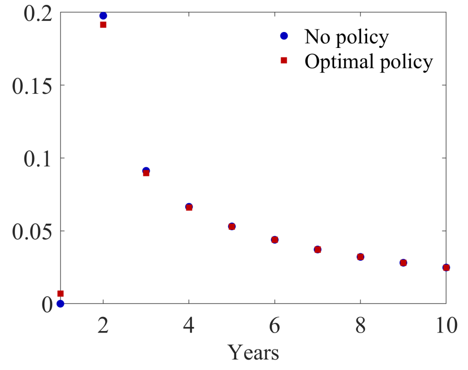

Figure 3.12: Mean MPC, for optimal temporary loan-to-value policy, over time

Note: Consumption responses under a temporary stricter LTV policy of 86 percent in $t = 1$, that is reversed in $t = 2$. Unexpected wealth shock of $-4,000$ dollars in $t = 2$. The consumption responses are normalized by $-4,000$ dollars also in $t = 1$, where there is only a change in policy and no shock has occurred. In the baseline model with no temporary policy, the LTV limit is 90 percent.

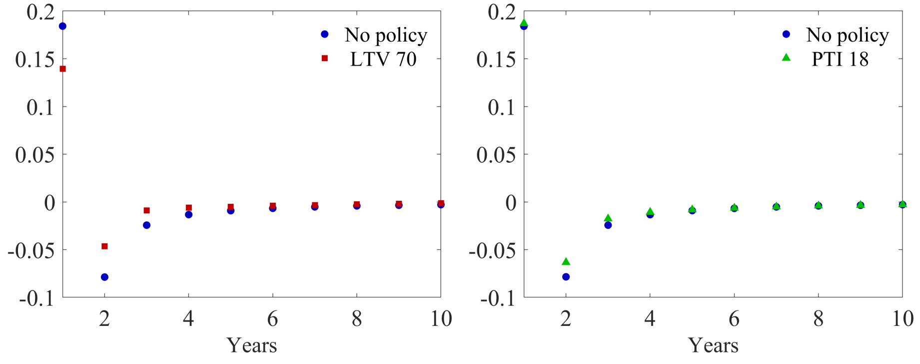

(a) Mean MPC over time **(b)** Mean MPC over time

Figure 3.13: MPCs for temporary LTV or PTI policy in boom-bust episode

Note: Consumption responses under a temporary stricter policy in $t = 1$, that is reversed in $t = 2$. Unexpected wealth shock of 4,000 dollars in $t = 1$ and $-4,000$ dollars in $t = 2$. The consumption responses are normalized by 4,000 dollars in all periods. In the baseline model with no temporary policy, the LTV limit is 90 percent and the PTI constraint is 28 percent.

Chapter 4

Inertia of dominated pension investments: evidence from an information intervention*†

*This study is registered in the AEA RCT Registry and the unique identifying number is: AEARCTR-0003139.

†This paper has been jointly written with Louise Lorentzon. We are grateful for guidance by Jonathan de Quidt, Per Pettersson-Lidbom, and Roine Vestman; and we are thankful for helpful discussions with Ingvild Almås, B. Douglas Bernheim, Luca Braghieri, Johannes Haushofer, Thomas Jansson, Paul Klein, Per Krusell, Andreas Madestam, Sean Myers, Arash Nekoei, Alessandra Peter, David Strömberg, Anna Tompsett, David Y. Yang, Robert Östling, and seminar participants at the Pension group at SNS, Stockholm University, Sweden's Financial Supervisory Authority, and the Swedish Pensions Agency. We are thankful to Montassar Ghachem for excellent web development services, and to the helpful staff at E-Print AB. We thank the Swedish Pensions Agency for sharing data. We gratefully acknowledge funding provided by Handelsbanken's Research Foundations, Torsten Söderbergs Stiftelse, the Royal Swedish Academy of Sciences, Stiftelsen Infina, and the Thule Foundation's research program on long-term savings. The authors have no relevant or material financial interests that relate to the research in this paper. All errors are our own.

4.1 Introduction

Although there are many mutual funds available, there is imperfect competition in the fund market. Fund companies make large profits, and there are high price disparities (Eriksson, 2014). With over 9 000 mutual funds from which to choose, an American household faces a rather challenging decision. When confronted with such a complex choice, many people select dominated alternatives resulting in lower consumer surplus and higher firm markups.[1][2] A relatively new strand of literature finds that differences in returns to savings across households account for a large share of wealth inequality; see Hubmer et al. (2020) for an overview. Additionally, there is a concern that people seldom reevaluate their investment choices.[3] A low level of active savers could impair the competition in the fund market by discouraging new entrants and thus allowing for higher markups among incumbent fund companies. This indicates that there are sizeable frictions affecting the demand side of the fund market: consumers' fund choices.

In this paper, we study why savers choose and stay with dominated funds. In a large-scale field experiment we examine three main hypotheses: lack of awareness of the price dispersion, search costs, and financial illiteracy. Further, we investigate if inertia in investments can be overcome by eliminating search costs for dominating alternatives. Knowledge about the causes of dominated choices can be an important input in the analysis of potential policies to counteract the imperfect competition in the fund market.

The experiment is conducted among savers in the Swedish Premium Pension system. The Premium Pension is a defined contribution part of the public pension where households are allowed to choose how their savings are invested. At the time of the experiment, there were over 800 funds from which to choose. In this set of funds, we identified two

---

[1]We define a fund to be dominated if the fund states that it is an index fund, and where there exists a cheaper index fund, in terms of fund fee, that follows the same index strategy. Hence, the funds are identical in terms of expected holdings. This is further described in Section 4.2.2.

[2]See, for example, Ayres and Curtis (2015) for a discussion of the problem with dominated funds in 401(k) plans; and Sinaiko and Hirth (2011) and Bhargava et al. (2017) for examinations of dominated choices in health insurance plans; and Johnson et al. (2015) regarding failures in mortgage refinancing.

[3]See, for example, the Swedish public inquiry by Billberg and Thomsson (2016).

dominated index funds, where there exists an identical fund in terms of expected holdings, with a lower fee. Our treatments provide information that increases the awareness of price dispersion, eliminates search costs, and aims at reducing financial illiteracy, to savers in the dominated funds.[4] Four different treatment letters are randomized to the savers, and in a subset of the letters, an incentivized search task to find the dominating fund is given.

Our results show that a lack of awareness that a selected fund is dominated, as well as search costs to identify the dominating alternative, contribute to dominated financial investment choices. Over a three-month period, 0.1 percent of our control group switches to the dominating fund. Informing savers that they have chosen a dominated fund increases this fraction by 1.9 percentage points. When reducing the search cost by providing the name of the dominating fund, the share increases by an additional 1.2 percentage points. However, providing an estimate of the expected future monetary gain from immediately reallocating the savings to the dominating fund does not influence the investment decision beyond what information on fee differences alone accomplishes. Thus, in contrast to previous studies, we do not find support for an exponential growth bias among the savers.[5]

Although the treatment effects regarding awareness and search costs are statistically and economically significant, one may still find the effects to be rather small, in particular for the treatment where search costs are essentially eliminated. One potential reason is a low take-up of the information. In fact, the confirmed reading share of our letters is 11 percent. This reading confirmation share can be used to approximate an upper bound of the treatment effect on the treated. Almost 30 percent of savers who are sent the treatment letter that provides the name of the dominating fund switch to this fund, when the treatment effect is scaled by the reading confirmation share.[6] This is indeed a much larger estimate.

[4]Salop and Stiglitz (1977) study search costs and show that the cost of gathering information alone can cause the equilibrium price to be non-competitive. Lusardi and Mitchell (2007) emphasize financial illiteracy; they discuss that many people are unfamiliar with basic economic concepts that are necessary for financial savings decisions.

[5]The exponential growth bias is found in, e.g., Levy and Tasoff (2016); Stango and Zinman (2009); Wagenaar and Sagaria (1975).

[6]Our estimated reading share can be considered as a lower bound, and thus the

Still, a majority of the savers do not switch to the dominating fund.

We conclude that providing information about comparable funds is a relatively cheap way to improve investments of many pension savers. As recent pension reforms in many countries often involve a shift in the responsibility from the government to the individual (see for example Ervik and Kildal (2016), Ring (2016), and Disney (2006)), the benefit of providing means to easily compare financial investments is likely to increase. However, there is still a large share of savers who do not respond to our information treatments. The high degree of inertia in pension investments, even when eliminating search frictions for identifying dominating alternatives, remains a puzzle.

There is a rich literature on inactivity and inertia in investments with a specific focus on the tendency to stay with default alternatives. Inertia is a feature that is inherently difficult to study in a laboratory, and therefore studies that aim to measure it are often observational. In particular, some papers have utilized changes in default contribution rates and investment allocations in 401(k) accounts. Typically, these studies find that there is a strong inertia in the default fund allocation (Madrian and Shea (2001), Choi et al. (2004), Choi et al. (2006)).[7] Madrian and Shea (2001) emphasize that the default investment allocation could be perceived as financial advice, and propose that procrastination caused by search costs, self-control problems, anchoring, and anticipated regret, could be contributing factors to the observed status-quo bias in investments. Regarding pension savers in general, Dahlquist and Martinez (2015) find a relatively large degree of inattention to past fund performance, as compared to other mutual fund investors. They suggest that savers may discount locked-in pension savings more than other savings.

Our sampling frame consists of investors who have opted out of the pension system's default fund in the past, and have a share of their portfolio in one of the dominated index funds. Thus, completely inert investors who never change their portfolios are excluded.[8] Hence, we

treatment effects scaled by this share provide upper bounds.

[7]Samuelson and Zeckhauser (1988) document a significant status quo bias in general, in a series of experiments; and Kempf and Ruenzi (2006) find that the status-quo bias is positively related to the number of available alternatives.

[8]A large share of savers in the Premium Pension system chose to opt out of the default fund at the inception of the system in 2000, when a “pro-choice” campaign was run, but have been inactive ever since (Cronqvist and Thaler (2004), Dahlquist et al.

can rule out suggested explanatory factors of inertia related to savings in default alternatives. In addition, since we study people who save in index funds where there exists a cheaper fund with the same index strategy, we can reject causes for inactivity related to differences in returns between alternatives, such as anchoring and anticipated regrets.

There is also an extensive literature focusing on the existence of dominated products in the financial industry. Hortaçsu and Syverson (2004) document that the most expensive S&P500 index fund had almost 30 times the fee of the cheapest fund in the category in 2000. Furthermore, approximately one third of all assets of U.S. all-equity mutual funds were managed by closet index funds in 2009 (Petajistoo, 2013). In these funds, the fund manager claims an active management, which justifies a higher fee, but the actual investments are close to those of comparable index funds. Carlin (2009) refers to studies that discuss the potential failure of the law of one price in many different markets: S&P500 index funds, money market funds, mutual funds, retail municipal bonds, credit cards, conventional fixed-rate mortgages, life annuities, and term life insurance.

On the theory side, Elton et al. (2004) discuss that in a market where there are uninformed investors and no arbitrage possibilities, the law of one price does not need to hold. Suggested explanations for the existence of dominated funds include: search frictions (1), information overload (2), lack of rationality among investors (3), costs charged by funds when savers alter investments (4), differences in the display of historical returns (5), and nonportfolio differentiation (6) (Elton et al. (2004), Hortaçsu and Syverson (2004), Agnew and Szykman (2005), Choi et al. (2010)).

In terms of the suggested reasons for dominated choices, we can rule out switching costs (4), differences in the display of historical returns (5), and most potential non-portfolio services (6). There are no switching costs in the Premium Pension system. In our sample of funds, the funds are ranked in the same order based on net returns, regardless of the historical period considered. Further, there are likely no significant non-portfolio services offered by the funds, since the fund choice is administrated by the Premium Pension system.[9] Among the remaining hypothezised causes

(2017)). However, in our sample, most of the savers made an active investment choice after the launch of the Premium Pension system, as seen in Figure 4.24.

[9] In addition, we can narrowly specify the other services the two dominated funds in our sample offer. In one of the two funds, the fund argues that it actively participates at

for dominated choices, we confirm the finding in Choi et al. (2010) that search costs have a limited impact on dominated choices, as a majority of the savers prevail with their dominated funds even when search costs are eliminated (1). Moreover, we do not find support for exponential growth bias to be a contributing factor for dominated choices. For some of our letters that provide more information, we see smaller treatment effects compared to letters that provide less information. This could indicate support for the information overload hypothesis (2) in a real investment choice setting.[10]

Overall, many of the previously proposed causes for both inertia and dominated choices can in our setting either be ruled out or fail to explain why the majority of the savers remain with the dominated funds. Some suggested explanations that would be beneficial to explore further in future work include limited rationality (3), mental accounting, and a lack of trust. Nonetheless, our finding that information letters can significantly improve investments for savers provides some optimism regarding the potential for public policies.

The remainder of the paper proceeds as follows. In Section 2 we describe the institutional setting, the sample, and the data. Section 3 presents the methodology including the treatments, the treatment assignment, and the hypotheses that are tested. The main results are presented in Section 4. In Section 5 we show some exploratory analysis, and Section 6 concludes the paper. In Appendix 4.A we discuss how the paper relates to our published pre-analysis plan.

4.2 Background and sample

4.2.1 The Premium Pension

The Swedish pension system was reformed in 1999. Citizens were given the opportunity to choose funds for part of their public pension savings: the Premium Pension. As of January 2018, the choice set is 846 funds

the shareholders' general meetings. We study heterogeneous treatment effects between the two funds in our sample and find, if anything, a stronger treatment effect for the savers in this fund, see Figure 4.5.

[10] When we added a possibility to confirm the name of the dominating fund, this resulted in fewer switches. See the comparison of treatment A and A-certain in Section 4.5.1.

(Pensionsmyndigheten, 2018b). Everyone who has ever had taxable labor income in Sweden and was born after 1937 is part of the Premium Pension system (Pensionsmyndigheten, 2018a). The employer pays 2.5 percent of each employee's labor income to that individual's Premium Pension account (Pensionsmyndigheten, 2018a).[11] As the contribution rate is specified in the system, a pension saver is only concerned with the investment decision. This enables us to study the fund choice in isolation from the inter-temporal savings decision. By the end of 2017, the fund volume managed in the Premium Pension system surpassed 1 137 billion SEK (approximately 125 billion USD) (Pensionsmyndigheten, 2018c).

4.2.2 Sample selection

We define a fund as dominated if it is an index fund that has a higher fund fee than another index fund that follows an equivalent index, i.e., the funds have the same investment strategy. In the Premium Pension system we identified two large and dominated index funds, denoted Fund^h (high fee) and Fund^m(medium fee). The dominating index fund is denoted Fund^l (low fee). Fund^m, Fund^l, and Fund^h were introduced in the Premium Pension system in 2000, 2006, and 2013, respectively. At all times, the funds have been ordered in the same way in terms of fund fees, as illustrated in Figure 4.1. Although all savers in Fund^m and Fund^h currently save in dominated funds, those who selected Fund^m before Fund^l was introduced (before 2006) did not make a dominated fund choice at the time of the choice. A description of the three funds and their historical performance is presented in Appendix 4.B.1.

The sample in this study consists of people aged 25–64 who save in either of the dominated funds, Fund^h or Fund^m, for their Premium Pension. For administrative purposes, the people in the sample have to be registered residents in Sweden. Individuals who save in both of the dominated funds, Fund^m and Fund^h, are dropped from the sample.[12]

[11] The contribution rate is capped when the labor income exceeds 7.5 income base amounts.

[12] Furthermore, we drop observations in case any of the main variables are missing: year of birth, labor income, or the fund choice information. In addition, we drop observations with missing or protected addresses since we are unable to treat these people (send them letters). A total of 43 letters were returned to us by the postal services. These individuals are also dropped from the analysis. In order to not keep

The final sample size is 29 662 people in the two funds combined, where about 60 percent save in Fund^m and 40 percent in Fund^h. Approximately five percent of the sample have allocated some Premium Pension savings to the dominating fund in addition to saving in one of the dominated funds. For a comparison of the individual characteristics of the savers in the two dominated funds with the Swedish population, see Appendix 4.B.2. We also present age and gender distributions and historical investment activity levels for savers who have all of their Premium Pension savings in the default fund, those who have opted out from the default, and savers in our two dominated funds; see Figure 4.14 and Figures 4.15 in Appendix 4.B.2. In addition, a comparison of the Premium Pension savings and shares allocated to the dominated and the dominating funds is provided in Table 4.7 in Appendix 4.B.2.

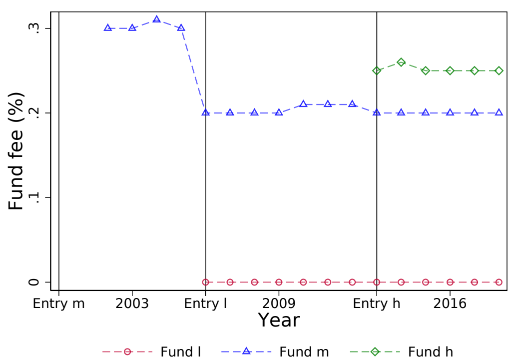

Figure 4.1: Historical fund fees
Note: Historical fund fees since the three funds entered the Premium Pension system. Fees denote net fees after the Premium Pension discount. The vertical lines marked by Entry m, Entry l, and Entry h, indicate when Fund^m, Fund^l, and Fund^h, respectively, were introduced in the Premium Pension system. The data is retrieved from the Swedish Pensions Agency in November 2018.

4.2.3 Data

In this project, we use data of fund characteristics, administrative data on individual level Premium Pension savings, investments, and background characteristics. The fund data on monthly historical gross returns and fees

observations with missing addresses in the control group, we exclude observations where the permanent residence municipality is missing.

(2000–2018) for the selected dominated and dominating funds is obtained from the Swedish Pensions Agency and from Morningstar Direct. From the Swedish Pensions Agency we also obtain data on individual Premium Pension fund choices and Premium Pension fund balances for the sample, before and after the experiment; as well as individual level data on labor income, gender, year of birth, marital status, and residential municipality; see Table 4.1, and Table 4.2.[13]

Variable	Range	Observation time	Data level
Savings share	[0,1]	Jun, Oct 2018	Individual
Savings amount (SEK)	$[0, \infty)$	Jun, Oct 2018	Individual
Investment date	[29 Sep 2000, 17 Oct 2018]	Jun, Oct 2018	Individual

Table 4.1: Premium Pension investment data
Note: Premium Pension investment data for the funds $Fund^l, Fund^m, Fund^h$. Savings share refers to the share of the Premium Pension savings that is allocated to each of the three funds $Fund^l, Fund^m, Fund^h$ (where 1 implies that all Premium Pension savings are saved in one fund). Investment date refers to the date of the most recent change of the Premium Pension portfolio. These variables are observed both before and after the information treatment. The data is retrieved from the Swedish Pensions Agency.

Variable	Range	Observation time	Data level
Annual labor income in 2016 (SEK)	$[0, \infty)$	2018	Individual
Year of birth	$[1953, 1993]$	2018	Individual
Gender	$\{Female, Male\}$	2018	Individual
Marital status	$\{S, M, D, W\}$	2018	Individual
Municipality	$\{1, 2, ..., 289, 290\}$	2018	Individual

Table 4.2: Individual background characteristics
Note: The table describes the data on individual level background characteristics. In 2018, the most recent available income data is for 2016. Annual income above the maximum annual PGI level (pension based income) is in 10 000 SEK intervals. The marital statuses S, M, D, W refer to single, married, divorced, widow/widower, respectively. Municipality refers to the municipality of registered residence. The data is retrieved from the Swedish Pensions Agency.

4.3 Experimental design

This section first describes the treatments, the treatment assignment, and the relevant outcome variables. We then present a theoretical framework

[13]An identification key of social security numbers is used to match names and addresses from the Swedish Tax Agency, when posting the letters.

of the decision to switch funds, followed by the hypotheses that are tested in the empirical analysis.

4.3.1 Treatments

We use four different treatment letters denoted: A, AN, AI, and ANI. In addition, there is a control group that receives no letter. A summary of the treatment letters follows and an illustration is provided in Figure 4.2.

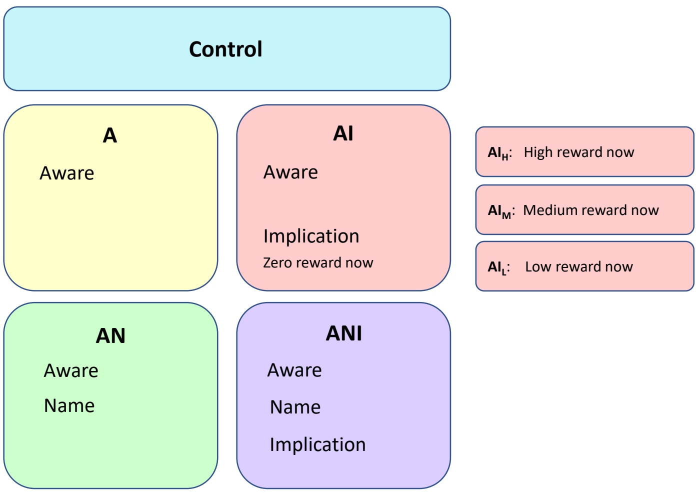

Figure 4.2: Illustration of the treatment groups

- **A** (Aware): A reminder of the name of the individual's current choice of fund (the dominated fund) and information that it is not the cheapest index fund with that investment strategy. The different fees of the current and the cheapest index fund with the selected investment strategy are stated. In addition, there is a short guide that describes how a fund selection is implemented and how one can categorize the current fund.
- **AN** (Aware+Name): The same content as letter A, plus information about the name of the cheapest index fund that follows the same investment strategy as the currently chosen index fund.

- **AI**$_a$ (Aware+Impl): The same content as letter A, plus a statement that clarifies the expected gain in the Premium Pension account balance at age 65, that they would get if they immediately switch their savings in the dominated fund to the dominating fund (the cheapest index fund in the chosen index fund category). This expected gain shows the monetary implication of the compounded effect of the fee difference. Furthermore, we include a search task with a range of immediate search rewards $a \in \{0, L, M, H\}$ if the name of the dominating fund is reported to us.[14]

- **ANI** (Aware+Name+Impl): The same content as letter AN, plus a statement that clarifies the expected gain in the Premium Pension account balance at age 65, that they would get if they immediately switch their savings in the dominated fund to the dominating fund (the cheapest index fund in the chosen index fund category).

The treatment letters are used to test three main hypothesis about *Awareness*, *Search costs*, and *Monetary implication*, as described in Section 4.3.5. A sample of treatment letter A (Aware) is found in Appendix 4.C.

Treatment compliance

We control that the treatment letters are sent to the savers. However, we do not observe whether individuals actually read the treatment letters. In order to test the compliance, we add a section to a randomly drawn sub-sample of letters in treatment A (Aware) and AN (Aware+Name), where the respondents are given an immediate reward if they confirm that they have read the letter.[15] An assessment of the share of recipients who read the letters enables us to estimate the size of the treatment effects conditional on receiving a treatment.

4.3.2 Treatment assignment

The treatment groups are randomly assigned within strata. We stratify based on covariates that we believe to be correlated with the outcome

[14]The low, medium, and high amounts correspond to 50, 250, and 700 SEK, respectively, which is approximately equivalent to 6, 28, and 77 USD. The savers can report the name of the dominating fund on a project website.

[15]The recipients are offered 200 SEK, which is equivalent to about 22 USD.

variables, in order to increase the estimation efficiency. The observations are divided into strata as follows, year of birth (two groups, split by the midpoint of the age range), labor income (two quantiles), fund choice (Fund^m, Fund^h), and fund share (two quantiles). If individual i belongs to strata s this is represented by the dummy variable $S_{is} = 1$. A fraction of the individuals who receive treatment A (Aware) and treatment AN (Aware+Name) are randomly assigned the reading task. Conditional on strata S_{is}, treatment k is independent of all other variables, including any potential outcomes.

Approximately 18 000 (17 960) individuals are sent treatment letters and roughly 4 800 individuals are left as controls.[16] The sample sizes by treatments are as follows.

- Control: 4 791
- A: 3 980, where 199 are sent the reading task
- AN: 3 986, where 199 are sent the reading task
- AI_0: 3 974
- $\text{AI}_{L,M,H}$: 2 006 (669, 669, 668 for each reward level, respectively)
- ANI: 3 991

An assessment of the treatment assignment is presented in Table 4.8 in Appendix 4.B.3, where pre-treatment characteristics across the treatment groups are compared. Figure 4.16 in Appendix 4.B.3 presents a comparison of the balance between the treatment groups and the control group in terms of the distributions of future rewards, savings in the funds, labor income, and age.

4.3.3 Outcome variables

The outcome variables for individual i, receiving treatment k, are observed three months after the treatment. The outcome variables of interest include:

- $Y_{ik}^{switch} \in \{0, 1\}$, an indicator variable that equals one if the individual switches from the dominated to the dominating fund;[17]

[16] A total of 23 people withdrew from the study, thus, the total sample size is 17 937.

[17] We define a switch as an observed reduction in the share of the Premium Pension portfolio invested in one of the dominated funds (Fund^m or Fund^h) and an increase in the portfolio share invested in the dominating fund (Fund^l).

- $Y_{ik}^{l} \in \{0, 1\}$, an indicator variable that equals one if the individual increases their portfolio share invested in the dominating fund, Fund^{l};
- $Y_{ik}^{mh} \in \{0, 1\}$, an indicator variable that equals one if the individual decreases their portfolio share invested in the dominated funds, Fund^{m} and Fund^{h};
- $Y_{ik}^{any} \in \{0, 1\}$, an indicator variable that equals one if the individual makes any investment change to their Premium Pension portfolio;
- $Y_{iAI_a}^{search} \in \{0, 1\}$, an indicator variable that equals one if the individual completes the search task (we study both those who successfully complete the task, and those who complete the task but provide an answer that is incorrect);
- $Y_{ik}^{read} \in \{0, 1\}$, an indicator variable that equals one if the individual confirms reading a letter.

The main outcome of interest is switches from the dominated to the dominating fund Y_{ik}^{switch} (hereafter referred to as *switch funds*). When we observe changes in the outcome variables we also observe the date at which a change was made. In addition, we observe logins at the Pensions Agency's website.[18]

Future reward from switching

We construct the variable *future reward from switching* R_{it}^{switch}, which is the additional amount in the Premium Pension account that individual i can expect to have at age 65, from immediately switching all of their savings in the dominated fund to the dominating fund.[19] This expected gain shows the monetary implication of the compounded effect of the fee difference. The number of years an individual has left until age 65 is denoted t.

The Swedish standard for pension forecasts is applied to construct the variable R_{it}^{switch}. The forecast depends on age, expected labor income, the fund balance in the dominated fund, the portfolio share in the dominated fund, the expected savings rate, the expected administrative fee of the

[18] We only have this data at the treatment group level.

[19] Age 65 was the mean age for starting to withdraw from the national public pension system, each year 2005–2016 (Carneck et al., 2017). Although there is no official pension age in Sweden, there is a norm for retiring at age 65, for example, the default retirement age in the public pension forecast is 65.

Swedish Pensions Agency, and the expected fund fees and returns of the funds. The computed forecast is rounded to the nearest 100 SEK, which is the practice for forecasts by the Swedish Pensions Agency. For a detailed description of how the forecast is computed and information about the distribution of forecasts, see Appendix 4.D.[20] We denote individual i's expectation of this future reward from an immediate switch by $E_i[R_{it}^{switch}]$.

4.3.4 Framework – the investment choice

In the setting of this study, awareness refers to being aware of that there exists a dominating fund to a currently chosen fund, i.e., that there exists a cheaper fund with the same expected holdings as a currently chosen fund. Our definition of search costs is the perceived cost of making a fund selection from the choice set in the Premium Pension system. For the main outcome variable (*switch funds*), the search cost is the effort to find the funds with the same investment strategy and to compare their fund fees. We refer to the effort of executing the choice, in the online Premium Pension account, as an administrative switch cost that is separate from the search cost. Moreover, we characterize financial illiteracy in this setting to be a lack of understanding of the expected monetary implications of changes in fund fees on the future Premium Pension account balance. This targets the documented tendency to underestimate the effect of compound interest.

To structure our thoughts about how the decision to switch from a dominated to a dominating fund is made, we consider the following framework. The choice depends on: individual i's discount factor of the future reward that is realized t years into the future β_{it}; the individual expectation of the future reward from immediately switching funds $E_i[R_{it}^{switch}]$; the size of any immediate search reward R_i^{search}; the individual perception of the immediate search cost associated with finding the dominating fund C_i^{search}; and the individual perception of the immediate cost associated with performing the switch C_i^{admin}.[21] Furthermore, we assume that the decision to switch from the dominated to the dominating fund requires

[20]One could consider this forecast to be a lower bound of the expected future reward associated with the information about the dominated and the dominating funds. Each individual could realize some additional benefits from the provided information, e.g., if the person adjusts other investments beyond the Premium Pension savings.

[21]We provide the immediate search rewards R_i^{search} in one treatment arm.

that an individual is aware of the fact that there exists a dominating fund to the currently chosen fund. We model this as a binary state captured by the indicator variable $I_i^{aware} \in \{0, 1\}$. The probability of being aware increases if receiving and reading one of our letters. However, being aware does not necessarily imply that a person will switch funds, as some people may not switch because they perceive the benefits to be too low compared to the costs. The outcome variable that equals one if a person switches funds is described by

$$Y_i^{switch} = f(\beta_{it}, E_i[R_{it}^{switch}], R_i^{search}, C_i^{admin}, C_i^{search}, I_i^{aware}) + \varepsilon_i.$$

The function $f(\cdot)$ is weakly increasing in $\beta_{it}, E_i[R_{it}^{switch}], R_i^{search}, I_i^{aware}$, and weakly decreasing in $C_i^{admin}, C_i^{search}$. Whether or not to switch from the dominated to the dominating fund can be modeled with the following decision rule,

$$Y_i^{switch} = I_i^{aware} \times 1[Z_{it} > 0], \tag{4.1}$$

where

$$Z_{it} = u(R_i^{search} - C_i^{admin} - C_i^{search}) + \beta_{it} u(E_i[R_{it}^{switch}]).$$

In our setting, reading treatment letter A (Aware) increases the probability of becoming aware that a dominating fund exists, I_i^{aware}. The immediate rewards that we provide in treatment AI_a (Aware+Impl) control the level of immediate search rewards R_i^{search}. Providing the name of the dominating fund, in treatment AN and ANI, eliminates the search cost of finding the dominating fund, C_i^{search}. Finally, the information about the expected monetary implication of a fund switch, in treatment AI and ANI, leads to an update of the expected gain from a switch, $E_i[R_{it}^{switch}]$.

4.3.5 Hypotheses

We are primarily interested in testing the following hypotheses.[22]

H1 *Awareness*
We test the null hypothesis that information about saving in a dominated fund does not increase the probability of a fund switch. This is tested

[22]Since treatments are randomized within strata, we condition on stratum dummies.

by comparing the probability of switching between treatment groups A (Aware) and Control (no letter):

$$H_0: \; E[Y^*_{i,A}|S_{is}] - E[Y^*_{i,Control}|S_{is}] = 0, \qquad (4.2)$$
$$H_1: \; E[Y^*_{i,A}|S_{is}] - E[Y^*_{i,Control}|S_{is}] > 0.$$

H2 *Search costs*
We test the null hypothesis that a reduced search cost, through information about the name of a dominating fund, does not increase the probability of a fund switch. This is tested by comparing the probability of switching between the treatment groups AN (Aware+Name) and A (Aware), as well as between the treatment groups ANI (Aware+Name+Impl) and AI_0 (Aware+Impl):

$$H_0: \; E[Y^*_{i,AN}|S_{is}] - E[Y^*_{i,A}|S_{is}] = 0, \qquad (4.3)$$
$$H_1: \; E[Y^*_{i,AN}|S_{is}] - E[Y^*_{i,A}|S_{is}] > 0,$$

and

$$H_0: \; E[Y^*_{i,ANI}|S_{is}] - E[Y^*_{i,AI_0}|S_{is}] = 0, \qquad (4.4)$$
$$H_1: \; E[Y^*_{i,ANI}|S_{is}] - E[Y^*_{i,AI_0}|S_{is}] > 0.$$

Both hypothesis tests for search cost, (4.3) and (4.4), identify the effect of including information about the name of the dominating fund. In contrast to the hypothesis test in (4.3), both information letters in the hypothesis test in (4.4) include the expected monetary implication of a switch. This difference enables us to test whether the effect of reducing search costs depends on if a saver is also informed about the expected monetary implication of a switch, i.e., a complementary effect.

H3 *Monetary implication*
We test the null hypothesis that information about the expected future monetary implication of a switch does not increase the probability of a fund switch. This is tested by comparing the probability of switching between the treatment groups AI_0 (Aware+Impl) and A (Aware), as well as between the treatment groups ANI (Aware+Name+Impl) and AN (Aware+Name):

$$
\begin{aligned}
H_0 &: \; E[Y^*_{i,AI_0}|S_{is}] - E[Y^*_{i,A}|S_{is}] = 0, \\
H_1 &: \; E[Y^*_{i,AI_0}|S_{is}] - E[Y^*_{i,A}|S_{is}] > 0,
\end{aligned}
\tag{4.5}
$$

and

$$
\begin{aligned}
H_0 &: \; E[Y^*_{i,ANI}|S_{is}] - E[Y^*_{i,AN}|S_{is}] = 0, \\
H_1 &: \; E[Y^*_{i,ANI}|S_{is}] - E[Y^*_{i,AN}|S_{is}] > 0.
\end{aligned}
\tag{4.6}
$$

Similar to the two hypothesis tests related to search costs, the difference between (4.5) and (4.6) allows us to test for complementary effects between reduced search costs and information about the expected monetary implication of a fund switch.

Implementation

All treatment letters are sent to the savers at the same point in time (July 19, 2018). The outcome variables Y^*_{ik} ($* \in \{switch, mh, l, any, search, read\}$) are observed three months after the treatment date, where the exact date of any change in the outcome variables is also noted. The search task and the reading task must be completed within a given time period (22 days) following the treatment date in order to receive the immediate rewards.[23]

4.4 Results

4.4.1 Estimation

Ordinary least square (OLS) regressions are used to estimate the treatment effects. Let T_{ik} be a dummy variable that takes the value one if individual i is sent treatment letter k, and zero otherwise. The main regression is given by

$$
Y^*_{isk} = \gamma_k T_{ik} + \delta_s S_{is} + \varepsilon_{isk}, \tag{4.7}
$$

[23]Participants complete the reading task and the search task on a project website, and the task rewards consist of general coupons that are valid in common stores in Sweden. The coupons are distributed to the agents at the project website.

where $k \in \{A, AN, AI_a, ANI\}$ and γ_k are the coefficients of interest. S_{is} are strata dummy variables, and ε_{isk} is an error term. Our primary tests use robust standard errors.[24] Unless otherwise stated, the data was retrieved from the Swedish Pensions Agency in October 2018, three months after the treatment date.[25]

4.4.2 Treatment effects

Main results

From regression (4.7) with the outcome variable Y^{switch}, we conclude the following. Information that makes the savers aware that there exists a cheaper index fund in the chosen category increases the probability of switching from the dominated to the dominating fund by 1.9 percentage points (0.1 percent makes the fund switch in the control group). Eliminating the search costs of finding the dominating fund, by providing its name, increases the probability of switching by an additional 1.2 percentage points. However, information about the expected monetary implication at retirement, of immediately switching to the cheapest fund, does not increase the probability of switching funds, beyond what information about the fee difference (in percent of fund balance) alone achieves. The results are illustrated in Figure 4.3 which shows the coefficient estimates with 95 percent confidence intervals. Table 4.3 displays the treatment effects with respect to our hypotheses listed in Section 4.3.5.

In addition to the outcome variable that states if a saver switches from the dominated to the dominating fund (Y^{switch}), we also investigate the outcome variables that indicate if a saver increases the share in the dominating fund (Y^{l}), decreases the share in the dominated funds (Y^{mh}), as well as any change in the fund portfolio (Y^{any}). The treatment effects on these outcome variables across treatment groups are documented in Table 4.4 and illustrated in Figure 4.4. The ordering of the treatment effects is similar for all outcome variables. The largest treatment effect is observed for the outcome variable that indicates a decrease of the share

[24] For robustness, we also calculate randomization inference p-values. The randomization inference p-values are simulated for other possible randomizations given the design, using Monte Carlo simulations. We use 100 repetitions, respecting strata.

[25] After three months there seem to be small differences in switching behavior between the treatment and the control group, see Figure 4.30.

Hypothesis	Treatments compared	Coef. diff.	P-value
H1 *Awareness*	A - Control	0.019	0.000
H2 *Search costs*	AN - A	0.012	0.001
	ANI - AI_0	0.020	0.000
H3 *Monetary implication*	AI_0 - A	-0.010	0.000
	ANI - AN	-0.003	0.476

Table 4.3: Hypothesis tests
Note: The outcome variable is the indicator variable for switching from the dominated to the dominating fund Y^{switch}. The equality in switches across treatment groups is assessed with Wald-tests. P-values from these tests are presented in column four.

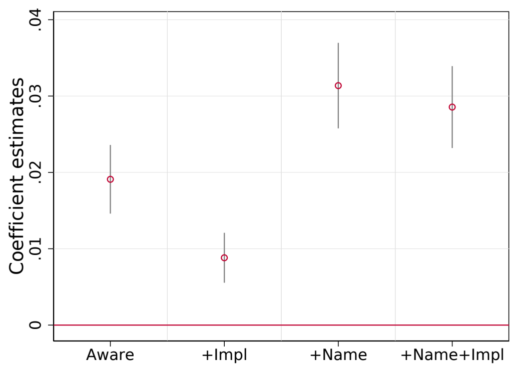

Figure 4.3: Treatment effects across treatment groups
Note: Regression coefficients γ_k from equation (4.7). The outcome variable is the indicator variable for switching from the dominated to the dominating fund Y^{switch}. The treatment effects are for treatments A, AI_0, AN, and ANI. The control group is the reference and has a mean of 0.001. The 95 percent confidence intervals are depicted around the coefficient estimates.

invested in one of the dominated funds. The probability of reducing the share invested in a dominated fund increases by 3.8 percentage points from making the savers aware of the fact that a dominating fund exists (Aware), and by an additional 1.1 percentage point when also providing the name of the dominating fund (Aware+Name).

To study the intensive margin of the treatment effects, we look at *how much* people change the portfolio share invested in the dominated funds. Further, we investigate the change in the portfolio share invested in the dominating fund for those who increase the share in this fund. The change in shares is measured relative to the initial share invested in Fund^m or Fund^h. In Figure 4.23 in Appendix 4.E, the distributions of these relative share changes are displayed. We note that most of the people who reduce the share invested in one of the dominated funds leave these funds completely. Among those who increase the share allocated to the dominating fund there is a larger dispersion in the relative share change, but the distribution is centered around 100 percent of what they had previously invested in the dominated fund. Appendix 4.E also presents the fractions of people that change their behavior in all treatment groups combined, compared to the control group, for different outcome variables.

	Outcomes			
Treatments	Switch	Increase L	Decrease M/H	Any change
Aware	0.019***	0.020***	0.038***	0.036***
	(0.002)	(0.002)	(0.004)	(0.005)
+Impl	0.009***	0.009***	0.023***	0.021***
	(0.002)	(0.002)	(0.004)	(0.005)
+Impl+Task Reward	0.018***	0.018***	0.033***	0.030***
	(0.003)	(0.003)	(0.005)	(0.006)
+Name	0.031***	0.032***	0.049***	0.046***
	(0.003)	(0.003)	(0.004)	(0.005)
+Name+Impl	0.029***	0.029***	0.040***	0.038***
	(0.003)	(0.003)	(0.004)	(0.005)
Observations	22,728	22,728	22,728	22,728
R^2	0.011	0.011	0.012	0.010

Table 4.4: Treatment effects
Note: Regression coefficients γ_k from regression (4.7), separately for each outcome variable. The control group is the reference and has a mean of 0.001, 0.002, 0.018, and 0.036, for the respective outcome variables. Robust standard errors are in parentheses, and the corresponding p-values are denoted by: *** p<0.01, ** p<0.05, * p<0.1. P-values from randomization inference are < 0.000 in all regressions.

Awareness and search costs

Our findings show that both a lack of awareness that a selected mutual fund is dominated, and search costs to find a dominating alternative, contribute to dominated financial investments. We reject both the hypothesis related to awareness (H1) and search costs (H2) at a one percent significance level. It is likely the case that many of the savers in the dominated funds are not aware that there exists a cheaper fund with the same index strategy. Nonetheless, only two percent of the savers who are sent letter

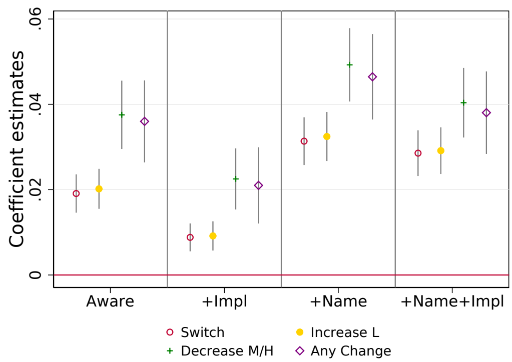

Figure 4.4: Treatment effects across treatment groups - several outcome variables

Note: Regression coefficients γ_k from regression (4.7), with $* \in \{switch, l, mh, any\}$. The control group is the reference and has a mean of 0.001. The markers show the point estimates and the vertical lines indicate the 95 percent confidence intervals.

A (Aware) switch funds. Search costs are a commonly suggested cause of dominated choices in the literature (Elton et al. (2004), Hortaçsu and Syverson (2004)). However, similar to Choi et al. (2010) we find that search costs alone cannot explain the high prevalence of savers prevailing with dominated funds. Indeed, the treatment effects from both the letter that increases awareness (A compared to Control) and the letters that eliminate search costs (AN and ANI compared to A and AI_0, respectively) are statistically and economically significant, but the effects are still quite small. This is particularly remarkable for the treatments where search costs to find a dominating alternative are eliminated.

Do people exhibit exponential growth bias?

One potential reason why relatively few savers switch funds when being sent letter A (Aware) and AN (Aware + Name) could be that they underestimate the expected future reward associated with a fund switch. There is a documented tendency that people underestimate the effect of compound interest, i.e., many people exhibit an exponential growth bias; see Goda et al. (2015), Levy and Tasoff (2016), and Stango and Zinman (2009). They find that the exponential growth bias reduces savings. In the monetary implication hypothesis (H3), we test if the exponential growth bias can also influence the investment decision by lowering the perceived future benefits of a lower fee. We compare the treatments where the expected future reward of switching funds is stated (ANI and AI_0), to the treatments where only the fund fee difference (in percent of fund balance) between the dominated and the dominating fund is presented (AN and A).

In Figure 4.3 and Table 4.3 we see that information about the expected monetary implication at retirement, of immediately switching to the cheapest fund, does not increase the probability of switching funds, beyond what information about the fee difference alone achieves. This result does not support previous findings that people tend to underestimate the effects of compound interest.

Our findings suggest that informing savers about the expected future gains from switching to a dominating fund, in addition to information about the fund fee difference, does not improve investment allocations. Although this result does not support an exponential growth bias among the savers, we cannot rule out its presence. When computing the forecasts for the expected future reward from switching we have to make a number of assumptions related to for example the future growth rate of labor income and the returns of the funds. It could be the case that the growth rate and the returns in our calculations are lower than what people expect, thereby counteracting a potential exponential growth bias.

In terms of the search cost hypothesis (H2) we note that the treatment effect is larger for the group that also receives the future monetary implication information (the coefficient difference between ANI and AI_0 is 0.020, and the coefficient difference between AN and A is 0.012). We conclude that even though the information about the expected monetary implication of a switch alone does not increase the probability of switching funds, it has complementary effects with the treatment that eliminates the search costs. Interestingly, there is an asymmetry in the complementary effect. When the future monetary implication of a switch is known, additional information that eliminates search costs increases the likelihood of switching funds. However, if the name of the dominating fund is known, additional information about the future monetary implication does not have an impact on the probability of switching funds.

How many read our letters?

Of those who received the reading task, 11 percent confirmed reading the letter.[26] Our main results are presented as intention-to-treat effects, i.e., the effect on people who were sent a letter. To approximate the treatment effect on the treated, i.e., for those who read the letter, we scale the intention-to-treat effects by the reading confirmation share of 11 percent.[27] Clearly, using the estimate for the share that read our letters to scale the treatment effects only yields an approximation of the treatment-on-the-treated effects. Further, since the reading confirmation estimate can be viewed as a lower bound for the share that actually read the letters, the scaled estimates provide upper bounds for the treatment-on-the-treated effects.[28]

The upper bound effect of the awareness letter for our main outcome variable *switch funds* Y^{switch} is 17 percentage points (0.019/0.11). Notably, a majority of the savers do not respond to the treatment even after rescaling the treatment effect by the reading confirmation share. In particular, this also holds for treatments AN (Aware+Name) and ANI (Aware+Name+Impl), where the search frictions to identify the dominating fund are eliminated. The largest estimate for the treatment-on-the-treated effect is obtained from treatment AN (Aware+Name), for outcome variable Y^{mh} (decrease in $Fund^m, Fund^h$). Here we find that the probability that a saver decreases the share invested in a dominated fund is at most 45 percentage points higher than for the control group. This is indeed a large estimate but nonetheless, it is puzzling that most of the savers remain with the dominated funds when they receive this treatment.

[26] 13.5 percent confirmed reading the letter in treatment A, and 8.5 percent confirmed reading the letter in treatment AN. These point estimates are not significantly different from each other. Further, we find no differences in characteristics between those who confirmed reading the letter as compared to those who did not confirm reading the letter, as displayed in Table 4.10 in Appendix 4.E.

[27] The treatment effect on the treated can be relevant for policy makers that have more channels through which they can reach out to citizens.

[28] The Swedish Pensions Agency estimates that 57 percent of people read their pension forecast letters in 2018 (Pensionsmyndigheten, 2019). The Pensions Agency finds that 83 percent received the letter, 78 percent of those who received it opened it, and 88 percent of those who opened it read it.

Heterogeneous treatment effects

In this section we explore how treatment effects vary with covariates. We are mainly interested in how treatment effects may differ by gender, age, labor income, and future reward (R_{it}^{switch}). For the continuous covariates, we primarily estimate heterogeneity across two quantiles. We also test for a heterogeneous treatment effect across the degree of urbanization of residential region and between savers in the two dominated funds (Fund^m and Fund^h). Let $\mathbf{X}_{ij}$ denote covariate j of which we investigate heterogeneous treatment effects. The heterogeneous treatment effects are tested using the following regression, estimated with OLS, where we interact the treatment dummies with the covariates $\mathbf{X}_{ij}$,

$$Y_{isk}^{switch} = \gamma_k T_{ik} + \eta_{kj} T_{ik} \mathbf{X}_{ij} + \rho_l \mathbf{X}_{ij} + \delta_s S_{is} + \varepsilon_{isk}, \tag{4.8}$$

where $k \in \{A, AN, AI_0, ANI\}$. The coefficients of interest are η_{kj}.

People with lower labor income and those who reside in relatively rural regions respond significantly less to almost all treatments. People with a lower income have previously been found to be less financially literate (Lusardi and Mitchell, 2007).[29] We also see that people who save in Fund^m react more than savers in Fund^h.[30] There are no heterogeneous treatment effects across gender, age, or the expected future reward from switching funds. The estimates from the heterogeneous treatment effect regression, for our main outcome variable switch of funds Y^{switch}, are displayed in Figure 4.5. The results look similar when we investigate heterogeneous treatment effects for the outcome variable Y^l. The effects are less pronounced and are in general not significantly different from zero for the outcome variables Y^{mh} and Y^{any}. In Figure 4.6, the heterogeneous treatment effect for labor income is illustrated over quintiles of labor income. It appears that the heterogeneous treatment effect stems from the whole labor income distribution. See Appendix 4.E.1 for additional analysis of heterogeneous treatment effects.

---

[29] Lusardi and Mitchell (2007) and Lusardi and Mitchell (2011) also find that more educated individuals are more financially literate. Since we do not have any information on educational background, we consider labor income to be a proxy for education.

[30] This also holds when restricting the sample to savers who made their fund choice after the introduction of Fund^l, i.e., when only considering individuals who made a dominated choice at the time of the choice.

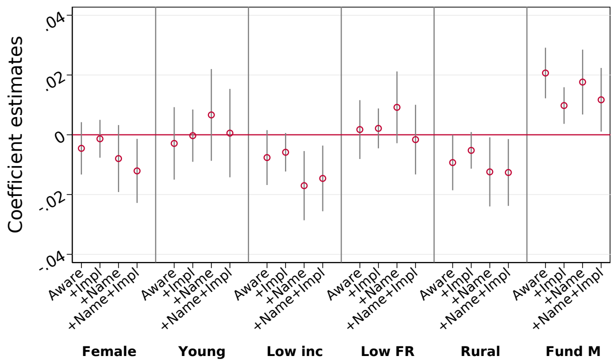

Figure 4.5: Heterogeneous treatment effects
Note: Heterogeneous treatment effects across gender, age, labor income, the expected future reward from a switch, residential region, and fund. Age, labor income, and the expected future reward are each separated into two quantiles; gender is divided into male and female, residential region is divided into urban and rural, and Fund^m is shown relative to Fund^h. The coefficient estimates show the relative differences in treatment effects to the other partition of the sample. The outcome variable is the indicator variable for switching from the dominated to the dominating fund Y^{switch}. The 95 percent confidence intervals are depicted around the coefficient estimates.

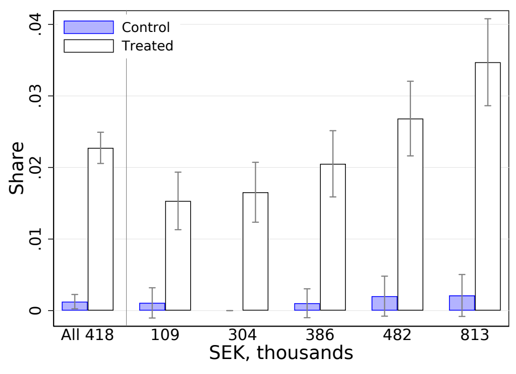

Figure 4.6: Treatment effects across labor income
Note: Shares who switch from the dominated to the dominating fund in all treatment groups combined versus the control group, across labor income quintiles. The vertical lines indicate the 95 percent confidence intervals. Mean earnings in each quintile are depicted under the corresponding bars. 1 SEK is approximately 0.11 USD. The two bars furthest to the left show the average shares.

4.4.3 Welfare effects of the information letters

In this section, we briefly discuss the cost of the information intervention and the expected gains stemming from the lower fees paid by savers who switch funds. Based on exponential discounting with an annual discount rate of 10 percent, the total discounted future reward that is saved by the people who switched all or parts of their savings from the dominated to the dominating fund amounts to 287 000 SEK.[31] The distribution of the expected discounted future rewards is displayed in Figure 4.7. The savings can be compared to the cost of producing and sending the letters. If we use a conservative estimate of 10 SEK per letter for printing and postal services, the cost of the treatment letters amounts to 180 000 SEK. Thus, even though the treatment effects may appear small in magnitude, an information intervention of this kind can still be motivated from a consumer cost-benefit analysis.

Overall, the fees that can be saved from switches from dominated to dominating funds are substantial. Choi et al. (2010) state that in 2007 over 200 million USD could have been saved in expenses on S&P 500 index funds alone, if savers had chosen the cheapest fund. One can also include a less strict definition of dominated funds by including closet index funds that tend to charge a higher fee than comparable index funds. Furthermore, additional savings may occur by informing people about the gains from switching from dominated funds, through the potential increased competition in the fund market, if funds respond by decreasing their fees. This effect would have an impact on all savers, not only those who switch funds. Thus, we argue that there are welfare gains to be made for consumers by providing information that facilitates the comparison of similar funds.

[31] In the control group, only 0.13 percent switched funds, compared to 2.3 percent in the treatment groups. We adjust for the fees saved by the control group. The total discounted future reward that would be saved if all people in the sample switched all of their savings in the dominated funds to the dominating fund, all else equal, would be almost 18 million SEK. 1 SEK is approximately 0.11 USD.

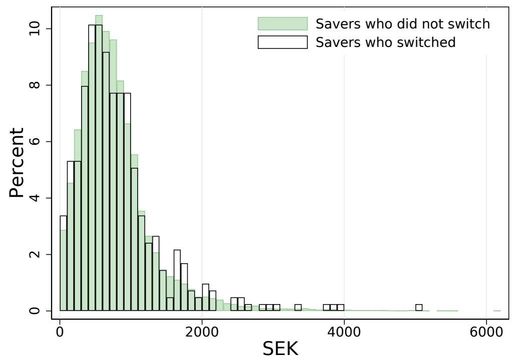

Figure 4.7: Discounted expected future rewards
Note: A distribution of the expected discounted future rewards, separately for those who did and did not switch from the dominated fund to the dominating fund. 1 SEK is approximately 0.11 USD. In total, 414 people switched and 22 314 people did not switch funds.

4.5 Exploratory analysis

In this section, we describe additional analyses regarding uncertainty, time preferences, and search costs. All these factors potentially influence the decision of whether or not to look for improvements in investments. We investigate distributions of search costs by using the variation in immediate search rewards R_{i0}^{search} and expected future rewards from switching R_{it}^{switch}.[32]

4.5.1 Does search uncertainty matter?

A saver that attempts to find a dominating fund may be uncertain regarding whether he or she has actually found a dominating fund.[33] To test the importance of this uncertainty for fund choices, we include an additional sub-treatment to treatment letter A (Aware). A subset of treatment letters, denoted A-certain, includes a one-attempt offer to verify the name of the dominating fund of the selected investment strategy, at the project website.

[32]See Appendix 4.E.2 for a brief discussion about potential procrastination and the duration of the treatment effects.

[33]Providing the name of the dominating fund eliminates this uncertainty. However, providing names of dominating funds may not be policy applicable.

By comparing treatment Aware and Aware-certain, we test for the importance of the uncertainty with respect to knowing if a fund is the dominating fund, for investment decisions. Our results indicate that this additional benefit does not significantly improve investment allocations. Rather, we see significantly smaller treatment effects when the option to verify the name of the dominating fund is given, see Figure 4.8 and Table 4.12 in Appendix 4.E. This may provide support for the information overload hypothesis in, e.g., Agnew and Szykman (2005). Another reason may be that the opportunity to verify the dominating fund signals that finding it may be difficult.

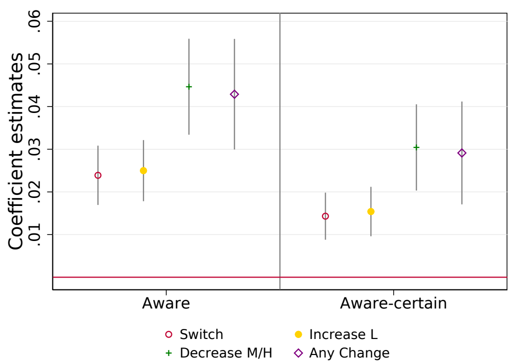

Figure 4.8: Treatment effects - uncertainty

Note: Regression coefficients γ_k from regression (4.7), with $* \in \{switch, l, mh, any\}$, reporting $k \in$ {A, A-certain}. The control group is the reference and has a mean of 0.001. The markers show the point estimates and the lines indicate the 95 percent confidence intervals.

4.5.2 Is a heavy discounting of pension savings the main cause of inactivity?

Dahlquist and Martinez (2015) argue that savers may discount locked-in pension savings more than other savings, and that people may use a separate mental account for retirement savings. In this section, we test the null hypothesis that an immediate compensation for the search cost does not increase the probability of switching funds. This is tested by comparing the probability of fund switches between different levels of search task rewards:

$$H_0: \; E[Y^*_{i,AI_a}|S_{is}] - E[Y^*_{i,AI_0}|S_{is}] = 0, \tag{4.9}$$
$$H_1: \; E[Y^*_{i,AI_a}|S_{is}] - E[Y^*_{i,AI_0}|S_{is}] > 0,$$

where $a \in \{L, M, H\}$.

We find that compensating for the search cost of finding the dominating fund increases the probability of a correct search, as shown in Figure 4.9. However, the larger share that correctly identifies the dominating fund does not directly translate into fund switches, as illustrated in Figure 4.10, and Table 4.11 in Appendix 4.E. The search rewards increase the probability that a saver searches and switches from the dominated to the dominating fund in general, but for the higher search rewards, there appears to be a greater discrepancy between the fraction of people who correctly identify the dominating fund and the share who actually switches funds.

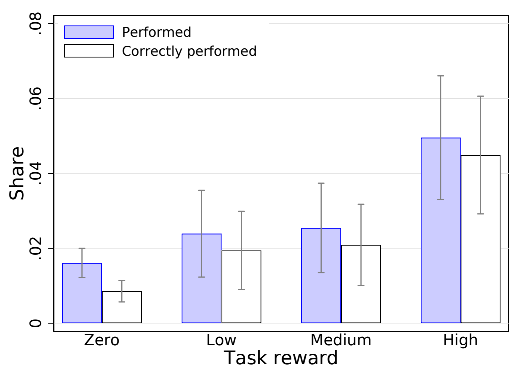

Figure 4.9: Completing the search task
Note: Shares that performed the search task and shares that correctly performed the search task, across search task reward levels, for treatment AI (Aware+Impl). The vertical lines indicate the 95 percent confidence intervals.

We are unable to exogenously vary the future reward levels of the savers and thus, it is difficult to compare how savers respond to increases in immediate rewards versus future rewards.[34] However, it appears as if

[34]For example, the expected future reward depends on age and labor income, which may be correlated with time preferences. The endogeneity problem in future rewards is further discussed in Section 4.5.3.

the given (partly endogenous) variation in future rewards across savers has a low correlation with the probability of switching funds, as discussed in Appendix 4.A. For inactivity to be explained by an extremely strong discounting, we would still expect to see that the probability of switching funds increases with the discounted future reward level in the ANI (Aware+Name+Impl) treatment, where the search costs are eliminated. However, this is not the case, as seen in Figure 4.11. That savers use a separate mental account for retirement savings, or exhibit a present bias, are left as possible explanations for why we observe increasing responses in immediate reward levels, but a low activity level across future reward levels.

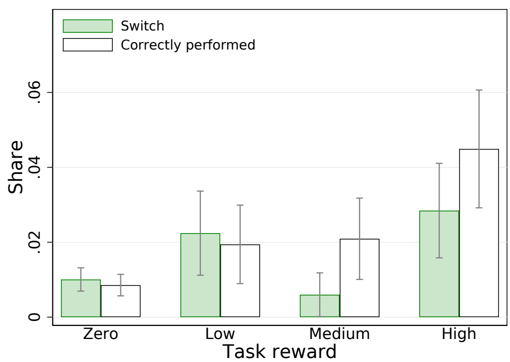

Figure 4.10: Search and switch
Note: Shares that correctly complete the search task and shares that switch funds, across search reward levels, for treatment AI (Aware+Impl). The vertical lines indicate the 95 percent confidence intervals.

4.5.3 How costly is the search for a dominating fund?

We are also interested in how costly people perceive a pension investment decision to be. To shed some light on this, we estimate a cumulative distribution of the search costs for finding a dominating fund. To estimate a search cost distribution, we make use of how the probability of switching funds varies with different immediate and discounted expected future rewards. For people who complete the search task and switch funds, we assume that the utility from the total discounted rewards (the immediate search reward and the discounted expected future reward from switching)

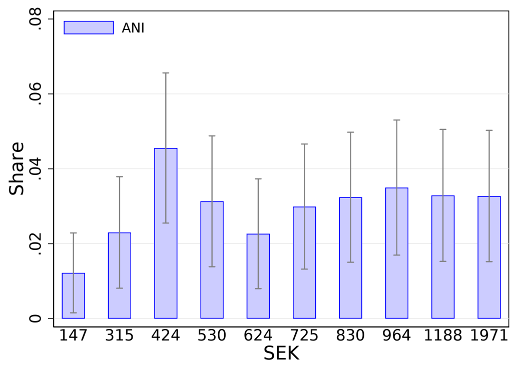

Figure 4.11: Switches across discounted future rewards
Note: Shares that switch funds, across deciles of discounted future rewards, for treatment ANI. The mean discounted future reward of each decile is depicted under the corresponding bar. 1 SEK is approximately 0.11 USD. The vertical lines indicate 95 percent confidence intervals.

exceeds the disutility of searching.[35] Thus, for a given immediate and discounted expected future reward, the estimated share of people switching funds provides an estimate for the share of people with search costs smaller than this amount.

In order to evaluate search costs *today*, we compute the discounted value of the expected future reward from switching for each individual. Dahlquist and Martinez (2015) suggest that people discount pension savings more than other savings. We also notice that many savers have large future rewards from switching funds (the mean expected future reward is 4 442 SEK), yet few people switch. We therefore choose a relatively low discount factor of 0.9 (yearly), and we use exponential discounting. This is a conservative choice, since it provides relatively low estimates of search costs.[36]

The different levels of immediate search rewards in treatment AI_a provide exogenous variation in rewards. Furthermore, we argue that some of the variation in future rewards is unrelated to search costs. The plausibly exogenous variation in expected future rewards stems from

[35]Since the expected future reward is stated in the AI_a (Aware+Impl) letters, we find it reasonable to assume that the individual expectation of the future reward is equal to this amount, i.e., $E_i[R_{it}^{switch}] = R_{it}^{switch}$.

[36]The shape of the estimated cumulative distribution is similar for different values of the discount factor.

differences in the fund fees of the dominated funds, the shares of the Premium Pension account balance allocated to the funds, the timing of historical investment choices, and the cap of contributions to the Premium Pension. We believe that these variations are not dependent on search costs and time preferences. However, there is still variation in future rewards that is likely correlated with search costs, e.g., through variation in labor income and age. Hence, we have to be cautious to interpret our estimated shares of switchers for different reward levels as a true cumulative distribution of search costs.

In Figure 4.12a, we see that only 18 percent of people are estimated to have search costs smaller than approximately 2 250 SEK.[37] Figure 4.12a presents the shares of savers switching from the dominated to the dominating fund over quantiles of total discounted rewards from switching. The figure displays the findings for treatment group AI (Aware+Impl), where immediate search rewards are varied. The shares are presented both in terms of the fraction of switches among those who were sent a letter, and scaled by the reading confirmation share. The scaled shares provide a more conservative assessment of the search costs, since it increases the share of households that can be assumed to have search costs lower than the values on the x-axis. Worth noting is that if we were to explain the low level of activity solely based on search costs, the cost for searching for the dominating fund appears very large. Roughly more than 80 percent of the savers find it more costly than approximately 2 250 SEK (about 250 USD) to search and switch to the dominating fund, as indicated by the rightmost bar with the scaled share. In Figure 4.12b, we also see that it is mainly the variation in immediate rewards that generates the differences in shares of switchers, i.e., there is no or little variation in search and switch behavior across levels of discounted future rewards.

[37] If we instead use a discount factor of 0.95 (0.85), 17 (23) percent of people are estimated to have search costs smaller than approximately 5 150 (1 300) SEK.

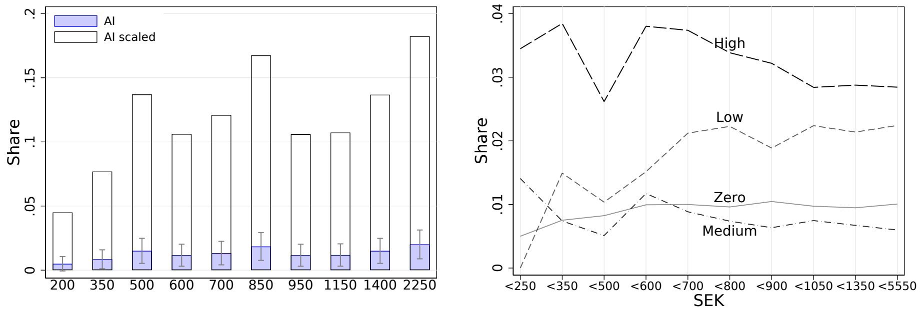

(a) Total discounted rewards (SEK) **(b)** Discounted future rewards (SEK)

Figure 4.12: Switches across discounted rewards

Note: (a) The share of people who switch funds, across deciles of total discounted rewards from switching (immediate and discounted expected future rewards), for treatment AI (Aware+Impl). The taller bars are scaled by the share who confirmed reading (0.11). The mean total discounted reward of each decile is depicted under each bar. 1 SEK is approximately 0.11 USD. The vertical lines indicate the 95 percent confidence intervals. (b) The share of people who switch funds, among savers with less than the indicated discounted expected future rewards on the x-axis (not including the immediate rewards). The shares are depicted separately for groups with different immediate rewards from switching, for treatment AI (Aware+Impl), not scaled by the share who confirmed reading. The numbers on the x-axis show the highest discounted future reward from switching in each decile of the sample.

4.6 Concluding remarks

To improve the competition in the fund market, it is crucial to understand why people choose and stay with dominated funds, and what information and choice architecture can support the decision making process of savers. Improving the choice environment by salience of relevant information can be a cheap way to increase consumer utility and competition.

In order to study causes for dominated fund choices, we run a large-scale field experiment in the Swedish Premium Pension system. Information letters are sent to people who save in dominated funds, where we test hypotheses regarding a lack of awareness of price dispersion, search costs, and financial illiteracy. People's real Premium Pension fund choices are observed and compared across treatment arms, following the letter treatments. By studying a real investment choice among savers that have actively opted out from the default fund and save in a dominated fund, we are able to analyze both reasons for dominated choices as well as the potential inertia in pension investments among relatively active pension savers.

Our results show that information about the existence of a dominating fund, and removed search costs for identifying this fund, significantly increase the probability that savers switch from the dominated to the dominating fund. Further, information that explains the expected future monetary implication of the fee difference, i.e., the effect of compound interest, does not increase the probability of switching. Overall, our findings show that providing relatively simple information that compares mutual funds has the potential to improve the investment allocations of savers.

Nonetheless, we find that an overwhelmingly large share of savers do not minimize fund fees, even when search costs are eliminated. Thus, a key question remains: what causes the high degree of inertia in pension investments among previously active investors? We propose the following potential reasons to be explored in future work: procrastination and self-control problems; lack of understanding the concept of index funds; lack of trust in information; and disbelief in the own ability to understand the provided information, and thus a low willingness to act upon it. Choi et al. (2010) find that people who fail to minimize fees among index funds often have a feeling that they are not optimizing their choice, indicating a low confidence regarding the fund choice. In future studies, it would be informative to survey the savers about their attitudes towards information letters. Furthermore, it would be useful to test if treatment effects depend on the sender identity.

References

Agnew, J. R. and L. R. Szykman (2005). Asset allocation and information overload: the influence of information display, asset choice, and investor experience. *Journal of Behavioral Finance 6*(2), 57–70.

Ayres, I. and Q. Curtis (2015). Beyond diversification: The pervasive problem of excessive fees and dominated funds in 401(k) plans. *Yale Law Journal 124*(5), 1476–1553.

Bhargava, S., G. Loewenstein, and J. Sydnor (2017). Choose to lose: health plan choices from a menu with dominated options. *The Quarterly Journal of Economics 132*(3), 1319–1372.

Billberg, L. and P. Thomsson (2016). Fokus premiepension; betänkande av 2016 års premiepensionsutredning. Investigation 61.

Carlin, B. I. (2009). Strategic price complexity in retail financial markets. *Journal of Financial Economics 91*(3), 278 – 287.

Carneck, A., H. Karlsson, and A. Carlsson (2017, May). Medelpensioneringsålder och utträdesålder, m.m. Report VER 2017-1, The Swedish Pensions Agency.

Choi, J., D. Laibson, B. Madrian, and A. Metrick (2004). For better or for worse: Default effects and 401(k) savings behavior. In D. A. Wise (Ed.), *Perspectives in the Economics of Aging*, pp. 81–121. University of Chicago Press.

Choi, J., D. Laibson, B. Madrian, and A. Metrick (2006). Saving for retirement on the path of least resistance. In E. McCaffrey and J. Slemrod (Eds.), *Behavioral Public Finance: Toward a New Agenda*, pp. 304–351. New York: Russell Sage Foundation.

Choi, J. J., D. Laibson, and B. C. Madrian (2010). Why does the law of one price fail? An experiment on index mutual funds. *The Review of Financial Studies 23*(4), 1405–1432.

Cronqvist, H. and R. H. Thaler (2004). Design choices in privatized socialsecurity systems: Learning from the swedish experience. *American Economic Review 94*(2), 424–428.

Dahlquist, M. and J. V. Martinez (2015). Investor inattention: A hidden cost of choice in pension plans? *European Financial Management 21*(1), 1–19.

Dahlquist, M., J. V. Martinez, and P. Söderlind (2017). Individual investor activity and performance. *The Review of Financial Studies 30*(3), 866–899.

Disney, R. (2006). Actuarial-based public pension systems. In G. L. Clark, A. H. Munnell, and J. M. Orszag (Eds.), *The Oxford Handbook of Pensions and Retirement Income*, pp. 268–289. Oxford University Press.

Elton, E. J., M. J. Gruber, and J. A. Busse (2004). Are investors rational? Choices among index funds. *The Journal of Finance 59*(1), 261–288.

Eriksson, H. (2014, June). Kassaklirr för fondbolagen. Newspaper article.

Ervik, R. and N. Kildal (2016). From collective to individual responsibility? Changing problem definitions of the welfare state. In R. Ervik, N. Kildal, and E. Nilssen (Eds.), *New Contractualism in European Welfare State Policies*. Routledge.

Goda, G. S., M. R. Levy, C. F. Manchester, A. Sojourner, and J. Tasoff (2015, August). The role of time preferences and exponential-growth bias in retirement savings. Working Paper 21482, National Bureau of Economic Research.

Hortaçsu, A. and C. Syverson (2004). Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P 500 index funds*. *The Quarterly Journal of Economics 119*(2), 403–456.

Hubmer, J., P. Krusell, and A. A. Smith (2020, May). Sources of U.S. wealth inequality: Past, present, and future. Working paper.

Johnson, E., S. Meier, and O. Toubia (2015). Money left on the kitchen table: Exploring sluggish mortgage refinancing using administrative data, surveys, and field experiments.

Kempf, A. and S. Ruenzi (2006). Status quo bias and the number of alternatives: An empirical illustration from the mutual fund industry. *The Journal of Behavioral Finance 7*(4), 204–213.

Levy, M. and J. Tasoff (2016, 06). Exponential-Growth Bias and Lifecycle Consumption. *Journal of the European Economic Association 14*(3), 545–583.

Lusardi, A. and O. S. Mitchell (2007). Financial literacy and retirement preparedness: Evidence and implications for financial education. *Business Economics 42*(1), 35–44.

Lusardi, A. and O. S. Mitchell (2011). Financial literacy around the world: an overview. *Journal of Pension Economics and Finance 10*(4), 497–508.

Madrian, B. C. and D. F. Shea (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior*. *The Quarterly Journal of Economics 116*(4), 1149–1187.

Pensionsmyndigheten (2018a). Förstå din pension. `https://www.pensionsmyndigheten.se/forsta-din-pension/sa-fungerar-pensionen/premiepension-en-del-av-den-allmanna-pensionen`. Retrieved: May, 2018.

Pensionsmyndigheten (2018b). Sök och jämför fonder. `https://www.pensionsmyndigheten.se/mina-tjanster/fondtorg/sok`. Retrieved: May 2018.

Pensionsmyndigheten (2018c). Statistik om premiepensionssparande. `https://www.pensionsmyndigheten.se/statistik-och-rapporter/statistik/statistik-for-premiepension`. Retrieved: May, 2018.

Pensionsmyndigheten (2019). 2018 Årsredovisning. `https://www.pensionsmyndigheten.se/om-pensionsmyndigheten/vart-uppdrag-och-organisation/arsredovisningar`. Retrieved: March, 2019.

Petajistoo, A. (2013). Active share and mutual fund performance. *Financial Analysts Journal 69*, 73 – 93.

Ring, P. J. (2016). The pension contract - a matter of obligations. In R. Ervik, N. Kildal, and E. Nilssen (Eds.), *New Contractualism in European Welfare State Policies*. Routledge.

Salop, S. and J. Stiglitz (1977). Bargains and ripoffs: a model of monopolistically competitive price dispersion. *The Review of Economic Studies 44*(3), 493–510.

Samuelson, W. and R. Zeckhauser (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty 1*, 7–59.

Sinaiko, A. D. and R. A. Hirth (2011). Consumers, health insurance and dominated choices. *Journal of Health Economics 30*(2), 450 – 457.

Stango, V. and J. Zinman (2009). Exponential growth bias and household finance. *The Journal of Finance 64*(6), 2807–2849.

Wagenaar, W. A. and S. D. Sagaria (71975). Misperception of exponential growth. *Perception & Psychophysics 18*(6), 416–422.

4.A Comments on the pre-analysis plan

This section contains information about differences between the published pre-analysis plan, and the conducted study and the analysis undertaken in this paper. Let us first highlight some changes in notation. The title of the paper is changed from: "The Choice of Pension Funds – An Information Experiment", in the pre-analysis plan. We also renamed the treatment letters to make them more closely connected to their contents.

Previous name	New name
Basic1	A (Aware)
Basic2	A-certain (Aware-certain)
A	AN (Aware+Name)
B	AI_a (Aware+Impl)
C	ANI (Aware+Name+Impl)

Further, we renamed the hypothesis *Asymmetric Information* to *Awareness* and the hypothesis *Financial Illiteracy* to *Monetary implication*. The link between the awareness treatment letter and the conceptual framework is also clarified in the paper. In addition, we have in the paper highlighted the relation between the monetary implication and exponential growth bias.

Implementation

The experiment was conducted in one round only.[38]

Data

We did not apply for additional individual data from Statistics Sweden, as we mentioned that we would consider.

Analysis

In the pre-analysis plan, we stated that we would estimate discount factors. Below, we outline the approach. Unfortunately, the analysis is

[38] We pre-specified the possibility of conducting a second round, in order to adjust the total sample size and to get an observation at a different point in time.

uninformative due to too little variation in the probability of switching funds across different levels of expected future rewards.

In addition to the pre-specified heterogeneity analysis, we analyze heterogeneity across levels of residential urbanization. We also include in the paper some calculations of consumer welfare, which were not pre-specified.

Do older people pay more attention to pension investments?

Is there age-dependent attention to pension investments? To attempt to answer this question, we estimate average implied discount factors for cohorts of people, with a different time t until retirement. The variation in the immediate search reward R^{search} and the future reward from switching R_t^{switch} is used in the following regression analysis, where we estimate different coefficients for different cohorts (implied by t),

$$Pr(Y^{switch} = 1)_i = \gamma_{0t} + \gamma_{1t} R_i^{search} + \gamma_{2t} R_{it}^{switch} + \gamma_{3t} \mathbf{X}_i + \varepsilon_i, \quad (4.10)$$

where $\mathbf{X}_i$ are individual covariates, and t denotes years left until retirement. How the probability of switching is affected by future and immediate rewards tells us about how rewards occurring at different points in time are valued, i.e., about time preferences.

The present value (PV) of a reward that is realized t years into the future is given by $\text{PV}(R_t^{switch}) = \beta_t R_t^{switch}$. From equation (4.10) we have a coefficient γ_{1t} that captures how the probability of switching funds varies with the immediate rewards (R^{search}), for a cohort with t years until retirement. We define the average implied discount factor $\bar{\beta}_t$ as the factor that equates the average responsiveness in switches from changes in future rewards to that from changes in immediate rewards. The above regression can then be written as

$$Pr(Y^{switch} = 1)_i = \gamma_{0t} + \gamma_{1t} R_i^{search} + \gamma_{1t} \bar{\beta}_t R_{it}^{switch} + \gamma_{3t} \mathbf{X}_i + \varepsilon_i.$$

From the estimation of equation (4.10), we infer the average implied discount factor, for the age group with t years until retirement, as follows,

$$\hat{\bar{\beta}}_t = \frac{\hat{\gamma}_{2t}}{\hat{\gamma}_{1t}}.$$

To handle the likely endogeneity problem of expected future rewards, we run instrumental variable regressions where we utilize the variation in future rewards that we argue to be exogenous of time preferences and search costs.[39]

Unfortunately, there is too little variation in the dependent variable, switching funds, across different levels of expected future rewards from switching. However, we see some signs of age-dependent attention when we examine heterogeneous treatment effects across age, see Appendix 4.E.1. Individuals who are closer to retirement tend to react more strongly to the treatments, but the differences are not statistically significant.

4.B Descriptive analysis of funds, sample, population, and treatments

4.B.1 The funds

Fund	Index name	No. savers	Fee	Correlation
Fund^m	OMXS30	31 281	0.20	0.9982
Fund^h	SIX30	19 462	0.25	0.9960
Fund^l	SIX30RX	75 858	0.00	1

Table 4.5: Characteristics of funds
Note: Characteristics of the dominated funds, Fund^m and Fund^h, and the dominating fund, Fund^l. The three indices all track the performance of the 30 most traded shares listed on the Stockholm Stock Exchange. The data is retrieved from the Swedish Pensions Agency. The number of savers only include Premium Pension savers. Fee refers to the net fee after Premium Pension rebates and is stated as a share of the savings in the fund. Correlation refers to the monthly historical correlation with Fund^l, since the funds became available for the Premium Pension.

[39]Specifically, we use the variation in the expected future reward due to differences in the fund fees of the dominated funds and the shares of the Premium Pension account balance allocated to the dominated funds.

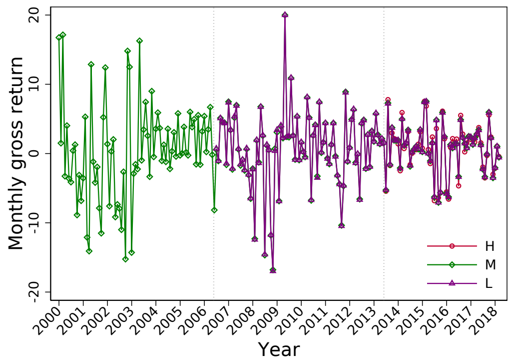

Figure 4.13: Historical monthly gross returns for the three funds

Note: The monthly gross return for the three funds Fund^l, Fund^m, and Fund^h. The gross return shows the return that investors would have received had they not paid any expenses. The gray vertical lines mark the entry of Fund^l and Fund^h, respectively. The data is from Morningstar Direct.

4.B.2 Sample and population comparison

	Sample				Population
Variable	Mean	Std. Dev.	Min	Max	Mean
Labor income (SEK)	416,718	376,614	0	$2.82e+07$	314,000
Year of birth	1967	7.6	1954	1993	
Female	0.46	0.50	0	1	0.50
Married	0.54	0.50	0	1	0.44
Future Reward	4,442	4,427	0	57,500	

Table 4.6: Characteristics of the sample vs the Swedish population

Note: Characteristics of the sample, savers in Fund^m and Fund^h, and a comparison with the Swedish population. The data for the savers in the two dominated funds is retrieved from the Swedish Pensions Agency in June 2018. Labor income refers to the annual labor income in 2016. The labor income in the population refers to the annual labor income in 2016 for the working age population in Sweden: age 20-64 (SCB). 1 SEK is approximately 0.11 USD. The fraction married in the population corresponds to people aged above 20 in Sweden in 2013 (Eurostat).

Variable	Mean	Std. Dev.	Min	Max
Savings Fund^m	49,383	60,323	0	654,782
Savings Fund^h	30,904	48,420	0	487,706
Savings Fund^l	70,839	47,140	0	389,969
Fund share, Fund^m	0.18	0.20	0	1
Fund share, Fund^h	0.10	0.16	0	1
Fund share, Fund^l	0.28	0.14	0	0.97

Table 4.7: Savings characteristics

Note: Savings (SEK) and fund shares for savers in Fund^m, Fund^h, and Fund^l. Savers in Fund^l are not part of the sample in this study, but are shown for comparison. 1 SEK is approximately 0.11 USD. The data is retrieved from the Swedish Pensions Agency in June 2018.

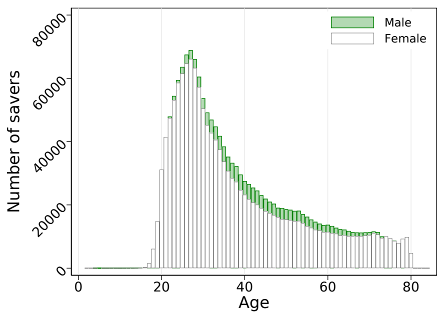

(a) All savings in the default fund

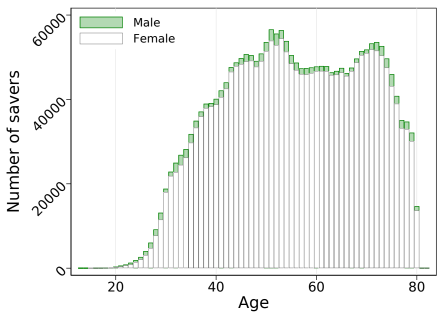

(b) Not all savings in the default fund

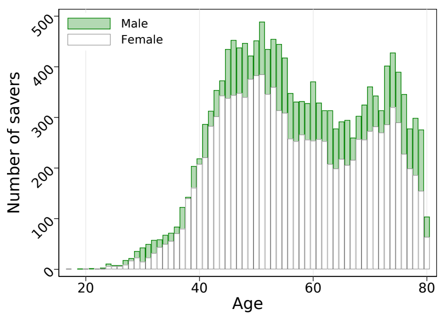

(c) Some savings in Fund^m

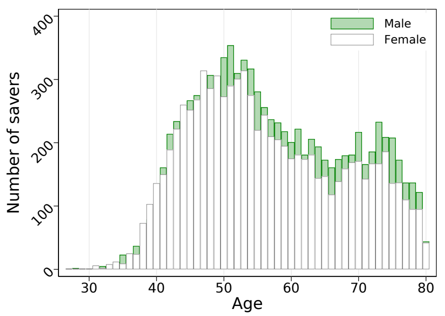

(d) Some savings in Fund^h

Figure 4.14: Age and gender distributions

Note: Age and gender distributions for all Premium Pension savers, displayed separately for those who save exclusively in the default fund, those who have actively chosen another fund than the default, and for savers in the two dominated funds. The data is retrieved from the Swedish Pensions Agency in November 2018, and shows the cross sections for June 18, 2018.

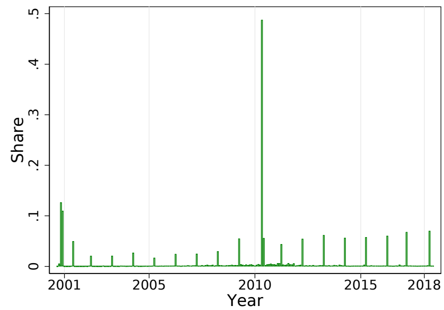

(a) All savings in the default fund

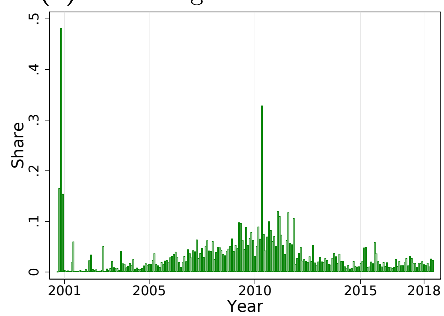

(b) Not all savings in the default fund

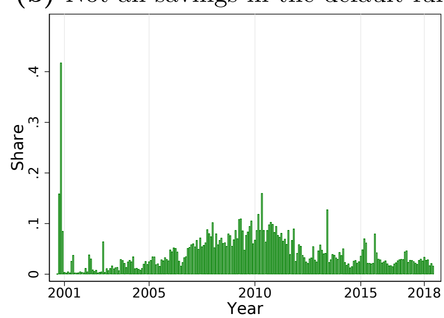

(c) Some savings in Fund^l, Fund^m, or Fund^h

Figure 4.15: Historical investment activity level

Note: Historical presentation of the fraction of savers who make an investment change for the Premium Pension savings, aggregated monthly and displayed separately for those who in June 18, 2018 save exclusively in the default fund, have actively chosen another fund than the default, and savers in Fund^l, Fund^m, or Fund^h. The periodic spikes that occur in December each year correspond to the inflow of new savings for the given year. If there is no change in the investment allocation, the new savings are allocated according to the previously chosen investment shares. In 2010, consultancy firms were allowed to perform automatic switches on behalf of Premium Pension savers, resulting in the significant spike of switches in that year. The data is retrieved from the Swedish Pensions Agency in November 2018.

4.B.3 Balance across treatments

Table 4.8 presents the mean values of pre-treatment characteristics across treatment groups. Distributions of future rewards, fund values in the dominated funds, labor income, and year of birth in the treatment as compared to the control group are displayed in Figure 4.16.

	Female	Year of birth	Married	Income	Share in Fund M/H	Years since last change
Aware	0.46	1967	0.56	421,635	0.28	6.3
+Name	0.47	1967	0.54	419,714	0.28	6.2
+Impl	0.46	1967	0.54	414,891	0.28	6.2
+Name+Impl	0.47	1967	0.54	418,511	0.28	6.1
Control	0.44	1967	0.55	415,116	0.28	6.3

Table 4.8: Pre-treatment characteristics across treatment groups
Note: The table shows mean values of covariates across treatment groups. The covariates Female and Married are in shares. Income refers to the mean annual labor income (in SEK) in 2016. 1 SEK is approximately 0.11 USD. Share in Fund M/H is the portfolio share saved in $Fund^m$ and $Fund^h$. The right-most column shows the average time in years since the most recent investment change. The data is retrieved from the Swedish Pensions Agency in June 2018.

	Control	Treatment	P-value
Annual income	415,116	418,265	0.62
Future reward	4,464	4,436	0.70
Savings in fund M/H	80,682	80,181	0.57
Year of birth	1967	1967	0.67
Share in fund M/H	0.28	0.28	0.80
Female	0.44	0.46	0.00
Married	0.55	0.54	0.61

Table 4.9: Balance of characteristics
Note: Mean values of pre-treatment characteristics across the control group, and the treatment groups combined. P-values are from t-tests of equality of means in the control group and the treatment groups combined. "Savings in fund M/H" is the value in the dominated fund. The variables Annual income, Future reward, and Savings in fund M/H are presented in SEK. 1 SEK is approximately 0.11 USD.

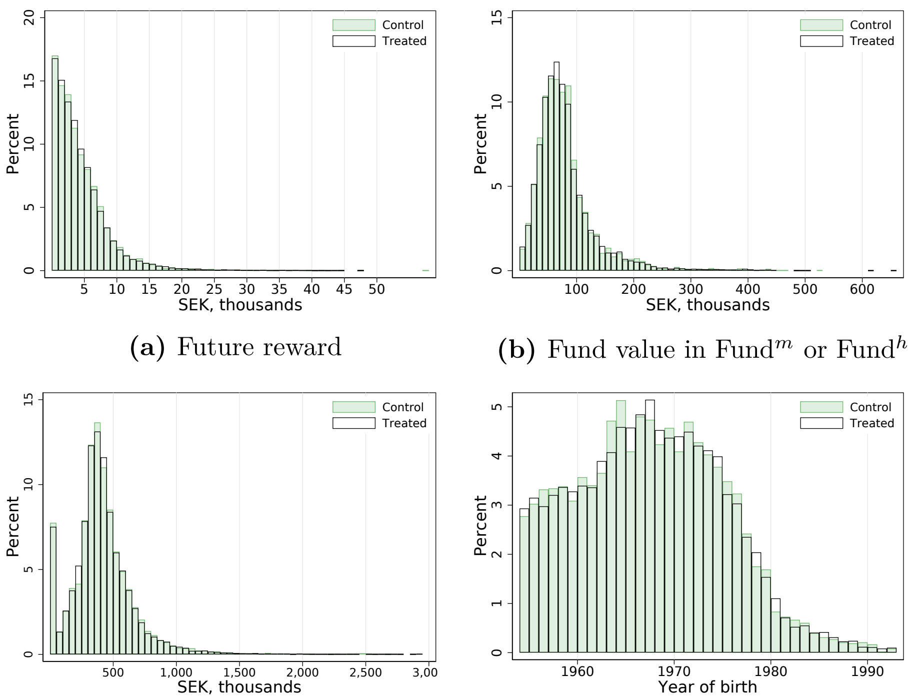

(a) Future reward

(b) Fund value in Fund^m or Fund^h

(c) Labor income

(d) Year of birth

Figure 4.16: Balance across treatment and control groups

Note: Distributions of savers in the treatment groups vs the control group, over (a) Expected future reward from a fund switch, (b) Fund value in the dominated fund, (c) Labor income, and (d) Year of birth. 1 SEK is approximately 0.11 USD. The data is retrieved from the Swedish Pensions Agency in June 2018.

4.C Information letter: Aware

Information om studien

Stockholms universitet

Det här brevet ingår i en akademisk studie som undersöker olika aspekter kring fondval. Syftet med studien är att ge ökad förståelse kring hur olika typer av information kan påverka fondval. Förhoppningen är att studien ska bidra till förbättrad information kring pensionsinvesteringar. Studien genomförs av forskare vid Stockholms universitet.

Personer som sparar sin premiepension i vissa utvalda fonder är med i denna studie. Vi som forskare använder pseudonymiserade uppgifter från Pensionsmyndigheten i vår analys. Eftersom uppgifterna är pseudonymiserade behandlar vi som forskare varken ditt personnummer, ditt namn eller din adress när vi analyserar uppgifterna. Dessa uppgifter används enbart för att administrera utskick. Uppgifterna som vi analyserar inkluderar fondval för premiepensionen, andel av premiepensionen som placeras i de utvalda fonderna, premiepensionssaldo i fonderna, datum då någon av fonderna valts, ålder, kön, civilstånd, inkomst samt bostadskommun. Den rättsliga grunden för behandling av personuppgifter är att forskning som bedrivs vid svenska lärosäten definieras som allmänt intresse. Vi behandlar din information så att obehöriga inte kan ta del av den. Studiens resultat kommer presenteras i en akademisk artikel. Resultaten och artikeln kommer inte att innehålla några personuppgifter. Den akademiska artikeln kommer att vara tillgänglig vid Stockholms universitet. Pseudonymiserade personuppgifter kopplade till studien bevaras hos forskarna tillsvidare, dock i minst tio år efter det att forskningsresultatet är färdigställt, i syfte att kunna göra fortsatta analyser samt för eventuell kontroll av studien.

Du har rätt att inge klagomål till tillsynsmyndigheten Datainspektionen/ Integritetsskyddsmyndigheten. Du har också enligt dataskyddsförordningen (EU 2016/679) rätt att i tillämpliga fall vända dig till Stockholms universitet och begära tillgång till personuppgifter som behandlas om dig och att begära rättelse eller radering av dessa personuppgifter i enlighet med artiklarna 15- 17. Du har rätt att begära begränsning av behandling av personuppgifter som rör dig (art 18) och att invända mot en sådan behandling (art 21). Du har även rätt att begära dataportabilitet i enlighet med artikel 20.

Du har härmed informerats om studien och har du frågor är du välkommen att kontakta oss.

Forskningshuvudman:	Stockholms universitet
Adress:	Universitetsvägen 10A SE-106 91 Stockholm
Tele:	08-16 20 00
Huvudansvarig forskare:	Per Petterson Lidbom
Personuppgiftsansvarig:	Stockholms universitet
Dataskyddsombud:	Benita Falenius
E-mail:	gdpr@su.se
Kontaktpersoner/forskare:	Karin Kinnerud och Louise Lorentzon
E-mail:	studie.pension@su.se

Figure 4.17: Letter, page 1

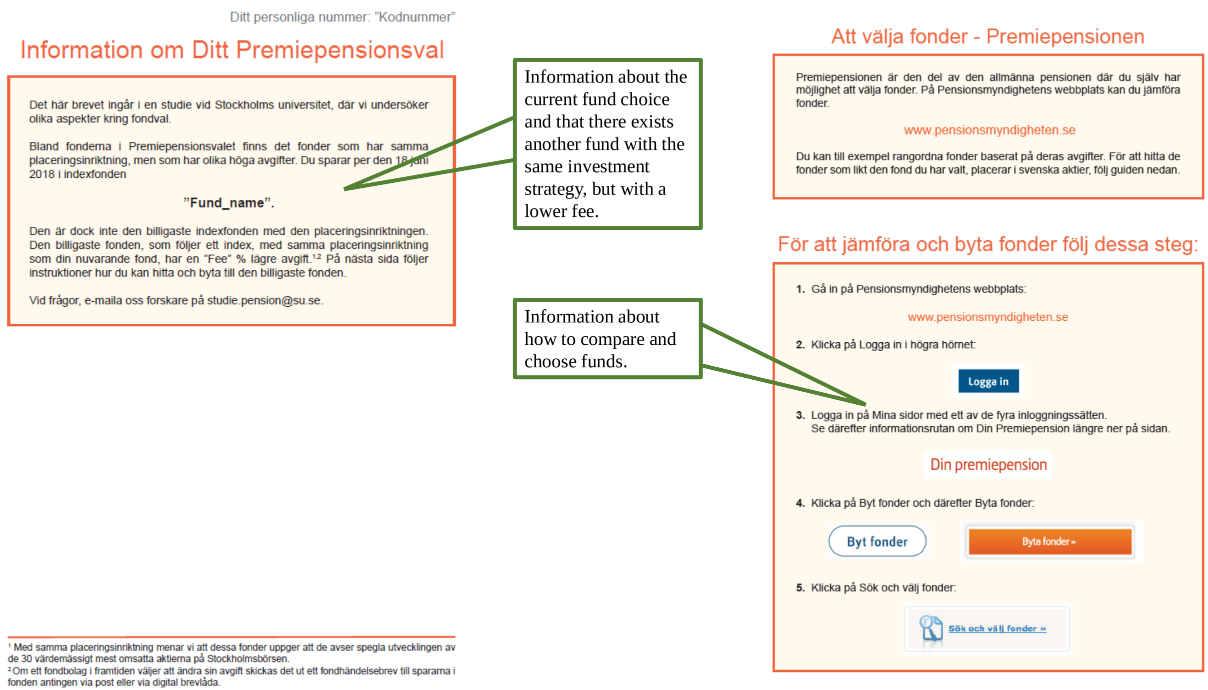

Figure 4.18: Letter, page 2 and 3

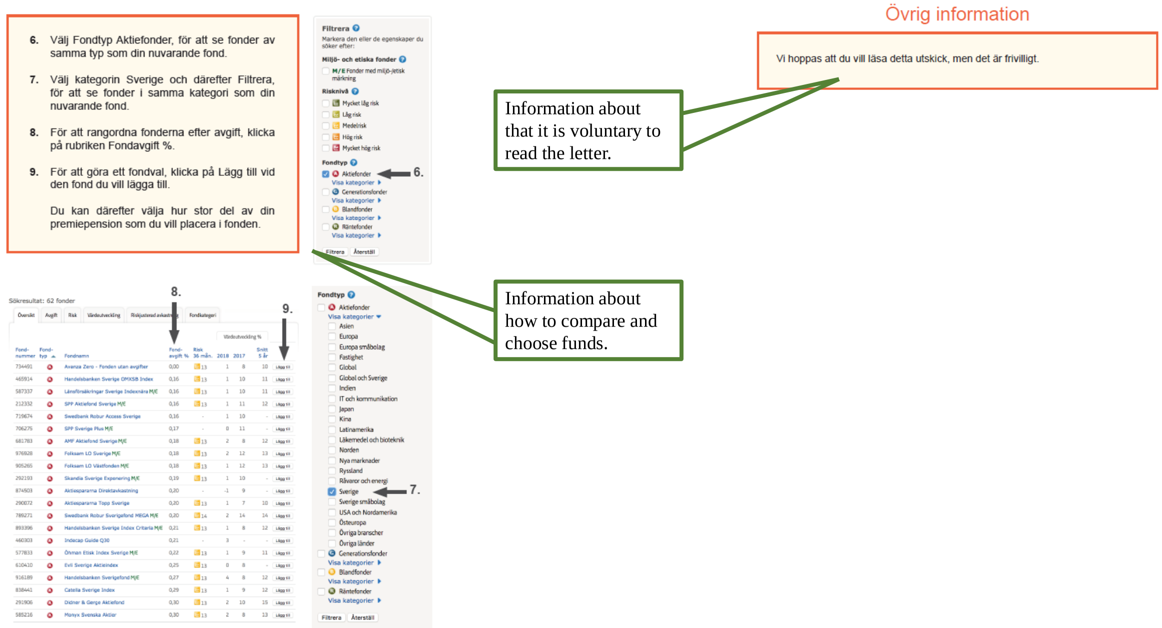

Figure 4.19: Letter, page 4 and 5

4.D Premium Pension forecasts and the expected future reward from a fund switch

In this section, the method and assumptions made to compute the expected future reward at age 65 from an immediate switch is described. We calculate the expected difference in pension savings, between saving in one of the dominated funds and saving in the dominating fund. First, we forecast individual Premium Pension account balances at age 65 under the two scenarios. Second, we compute the difference between the two forecasts. To compute the forecasts, we use individual level data on: age, labor income, fund balance in the dominated fund, portfolio share in the fund, savings rate, as well as data on the Swedish Pensions Agency's administrative fee, fund fees, and fund returns.

The forecasts are based on the following assumptions, from the standard for Swedish pension forecasts, where applicable. The frequency of timing is yearly.

Income

The yearly labor income for individual i, set at time t and paid at time $t+1$, is denoted y_{it}. The real labor income growth is assumed to be zero and thus, $y_{it} = y_i \, \forall t$.

Fund balance

The individual Premium Pension fund balance in one of the dominated funds, at time t, is denoted k_{it}. k_{it} corresponds to the fund balance at the beginning of year t, when the contribution from the previous year has just been added to the account.

For the forecasts, we assume that the portfolio shares in the chosen funds for the Premium Pension are constant over time. We also assume that the current fund balance in the dominated fund corresponds to the share of the Premium Pension account allocated to that particular fund.[40]

[40] If the chosen funds have realized different net returns since the allocation choice was made, this assumption may be violated. However, this assumption only has an impacts when the cap of the administrative fee is reached, and how much of the administrative fee is paid from the considered fund, as discussed below. The administrative fee is small relative to expected fund returns and fund fees, and hence its impact on the forecast is

Savings rate

The savings rate for the Premium Pension is a share s of labor income, up to an earnings cap y^{cap}. The savings rate and the cap are assumed to be constant for all future time periods. Thus, for individual i, at time t, the savings rate as a share of labor income is provided by

$$s_{it} = \min\{s, s \cdot y^{cap}/y_{it}\}.$$

Given the assumption of zero real labor income growth, the savings rate is constant over time for a given individual, i.e., $s_{it} = s_i \, \forall t$.

Administrative fee

The Swedish Pensions Agency's administrative fee is a share a of the total portfolio balance, up to an account balance cap AB^{cap}. The administrative fee is assumed to be constant for all future time periods. To compute if and for what time periods the administrative fee cap applies, we assume that the current share θ_i, of the Premium Pension account balance that is allocated to the dominated fund, remains the same in the future. In other words, we assume that the funds in the current portfolio have the same net return.[41] Thus, for individual i, at the end of year t, the administrative fee, as a share of the fund balance in the fund considered, is given by

$$a_{it} = \min\left\{a, a \cdot \frac{\theta_i \cdot AB^{cap}}{k_{it}(1+R)}\right\},$$

where R denotes the gross real rate of return of the fund.

Fund fees

The fund fee f is a yearly rate of the fund balance. We assume that the fund fees of the three funds are constant over time, $f \in \{f^h, f^m, f^l\}$.

---

low.

[41]The assumption of constant portfolio shares only affects when the cap of the administrative fee is reached, and how much of the administrative fee is paid from savings in the considered fund.

Fund return

The expected gross real rate of return of the funds is denoted R, and is assumed to be constant over time. The expected net real rate of return of a fund in year t is given by the gross real rate of return minus the fund fee and the administrative fee expressed as a share of the fund balance, i.e., $r_{it} = R - f - a_{it}$.

4.D.1 Premium Pension savings forecast computation

The expected fund balance, q years into the future, at year t, is given by

$$k_{i,t+q}(f) = k_{it} \prod_{n=0}^{q-1} (1 + r_{i,t+n}(f)) + \sum_{j=0}^{q-1} \left[s_i \theta_i y_i \prod_{n=q-j}^{q-1} (1 + r_{i,t+n}(f)) \right]$$

where

$$\begin{aligned}
&\prod_{n=q}^{q-1} (1 + r_{i,t+n}(f)) = 1 \\
&\prod_{n=q-1}^{q-1} (1 + r_{i,t+n}(f)) = (1 + r_{i,t+q-1}(f)) \\
&r_{it}(f) = R - f - a_{it} \\
&f \in \{f^h, f^m, f^l\}.
\end{aligned}$$

The current ($t = 0$) expected difference in pension savings, q years into the future, based on fund fee differences, is

$$\begin{aligned}
R_{iq}^{lh} &= k_{i,q}(f^l) - k_{i,q}(f^h), \\
R_{iq}^{lm} &= k_{i,q}(f^l) - k_{i,q}(f^m),
\end{aligned}$$

where R_{iq}^{lh} and R_{iq}^{lm} denote the differences generated by the high and medium fees compared to the low fee, respectively. These variables are what we refer to as R_{it}^{switch} in the paper.

4.D.2 Data

The individual labor income level is the most current yearly labor income available (2016). The individual fund balance in a dominated fund is retrieved in June, 2018. All new contributions are allocated in accordance with the reported portfolio shares as of June 2018.

A savings rate of $s = 2.5\%$ is based on the savings rate in 2018 and is assumed to be constant in all future time periods. With the 2016 labor income data, the savings rate cap applies when the earnings exceed $y^{cap} = 444\,750$ SEK (2016), based on

$$\begin{aligned} y^{cap} &= 7.5 \cdot 2016 \text{ income base amounts} \\ &= 7.5 \cdot 59\,300 \text{ SEK} \\ &= 444\,750 \text{ SEK}. \end{aligned}$$

The administrative fee in 2016 was $a = 0.11\%$ of the total portfolio balance up to a fee cap of 120 SEK. Thus, the fee cap applies when the total portfolio balance exceeds $AB^{cap} = 120/0.0011 = 109\,091$ SEK. For the yearly gross real rate of return of the funds, we use the Swedish Pensions Agency's assumption for stocks of $R = 2.7\%$.[42]

The forecast is stated in constant 2018 SEK. Data on labor income and the Pensions Agency's administrative fee cap are from the end of 2016. The letters with the forecasts were sent in July, 2018. We adjust the 2016 SEK data for inflation, using the Consumer Price Index (CPI) in January 2017 and May 2018 (the most current CPI level as of July 5, 2018). The inflation adjustment is given by

$$\pi_{adj} = \frac{I^{2018}_{May}}{I^{2017}_{Jan}} \approx 1.0325.$$

4.D.3 Distributions of expected future rewards

[42]The Swedish Pensions Agency's assumption of fund returns in general is $R = 2.1\%$, but since the funds we consider only contain stocks, we use their specific assumption for stocks.

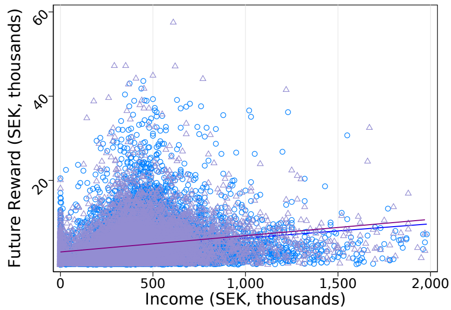

(a)

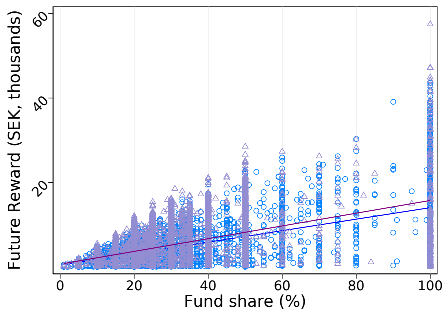

(b)

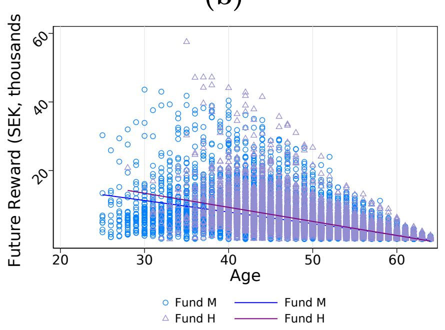

(c)

Figure 4.20: Expected future reward distributions

Note: Distributions of the expected future reward at age 65 from immediately switching from the dominated to the dominating fund, over current labor income (a), fund share allocated to the dominated fund (b), and age (c). The data is retrieved from the Swedish Pensions Agency in June 2018. Labor income in graph (a) is capped at 2 million SEK.

4.E Additional results and robustness checks

Unless otherwise stated, the data is retrieved from the Swedish Pensions Agency in October 2018. Figure 4.21 presents the shares of households, unconditional on strata, that change their behavior in terms of different outcome variables, for all treatment groups combined, compared to the control group.

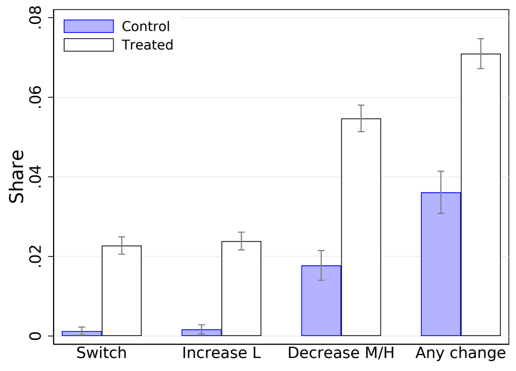

Figure 4.21: Outcomes across all treatment groups and the control

Note: Shares of changes across the different outcome variables (Y^{switch}, Y^{l}, Y^{mh}, Y^{any}), for all treatment groups combined compared to the control group. The lines indicate the 95 percent confidence intervals.

	Did not confirm	Did confirm	P-value
Annual income	410,207	429,654	0.47
Future reward	4,509	4,232	0.48
Savings in fund M/H	85,215	79,193	0.30
Year of birth	1967	1967	0.75
Share in fund M/H	0.28	0.27	0.52
Female	0.44	0.45	0.92
Married	0.58	0.63	0.31

Table 4.10: Characteristics across reading confirmation

Note: Mean values of characteristics for those who did not confirm, and those who did confirm reading the letter. P-values are from t-tests of equality of means. “Savings in fund M/H” is the value in the dominated fund. The variables Annual income, Future reward, and Savings in fund M/H, are presented in SEK. 1 SEK is approximately 0.11 USD.

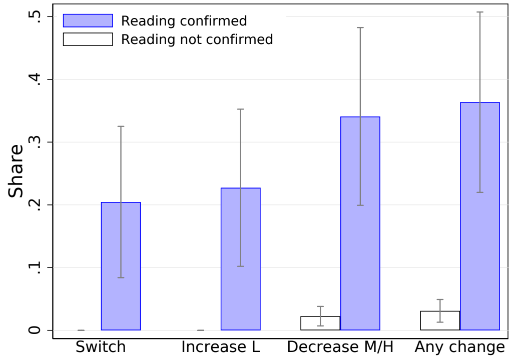

Figure 4.22: Outcomes for those who did and did not confirm reading

Note: Shares of changes for the different outcome variables (Y^{switch}, Y^{l}, Y^{mh}, Y^{any}), for those who received the reading confirmation task, across those who confirmed versus not confirmed reading the treatment letter. The lines indicate the 95 percent confidence intervals.

Reward comparison	Correct search		Switch	
	Coef. diff.	P-value	Coef. diff.	P-value
Low-Zero	0.011	0.05	0.019	0.04
Medium-Zero	0.012	0.03	-0.014	0.23
High-Zero	0.036	0.00	0.018	0.01

Table 4.11: Coefficient differences - search rewards

Note: The table shows the differences in coefficients across reward levels, separately for the outcomes correctly completing the search task, and for switching funds. The p-values are from Wald tests, testing the equality of the coefficients.

Figure 4.23 shows the changes in fund shares relative to the initial shares invested in Fund^m or Fund^h. The outcome variables for fund share changes are defined as follows,

- $Y_{ik}^{share,mh} \in [-1, 100]$, the change in the portfolio share invested in the dominated fund relative to the initial share;
- $Y_{ik}^{share,l} \in [-99, 100]$, the change in the portfolio share invested in the dominating fund, relative to the initial share invested in the dominated fund.[43]

[43]The fund share is chosen in increments of 0.01.

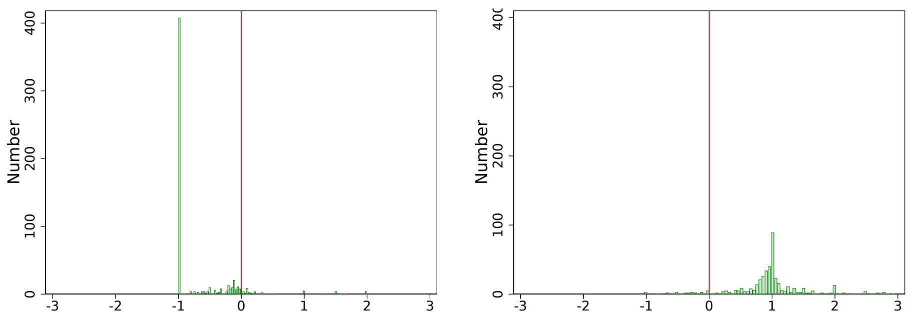

(a) Relative fund share changes: Fund^{mh}

(b) Relative fund share changes: Fund^{l}

Figure 4.23: Relative fund share changes

Note: Fund share changes relative to the initial share invested in Fund^{m} or Fund^{h}. The distributions show only those who made the relevant fund share change. For visual purposes, the x-axes are cut at -3 and +3.

	Aware	Aware-certain	P-values
Switch	0.024	0.014	0.032
Increase L	0.025	0.015	0.038
Any change	0.043	0.029	0.093
Decrease M/H	0.045	0.030	0.049

Table 4.12: Treatment effects - uncertainty

Note: Treatment effects for Aware and Aware-certain. The p-values are from Wald tests of the coefficient equality between the treatment groups Aware and Aware-certain.

4.E.1 Heterogeneous treatment effects

Levels of previous activity in the fund market

Figure 4.24 shows the distribution of the time of the most recent investment choice for the treated sample, and compares those who switched funds to those who did not. It can be noted that the savers who switched funds overall have made an investment change more recently. The distributions look similar for the other outcome variables (Y^l, Y^{mh}, Y^{any}). This finding is also apparent in Table 4.13, where we see that the average time since the most recent investment choice was more than a year longer among those who did not switch funds as compared to those who did. The shares that switch funds over the most recent investment date quintiles are displayed for the control and treatment groups in Figure 4.25.

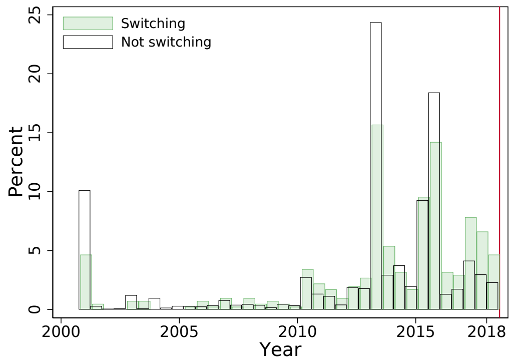

Figure 4.24: Time of most recent investment change
Note: Distribution of the time of the most recent investment change for the treated sample, separated into those who switched funds and those who did not switch funds. The red vertical line shows the treatment date.

	Control	Treated
No switch	6.27	6.23
Switch	4.84	5.13

Table 4.13: Years since most recent investment change
Note: Means of years since the most recent investment change before treatment, for the control group, and all treatment groups combined, separately for those who switched and those who did not. The data is retrieved from the Swedish Pensions Agency in October 2018.

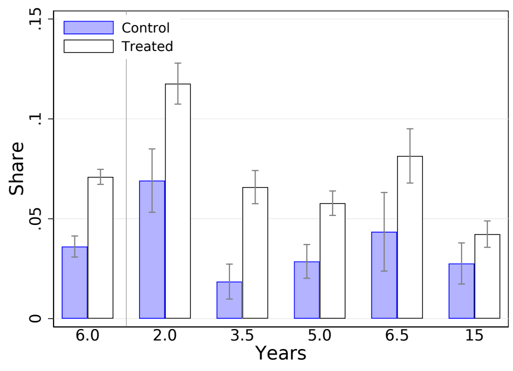

Figure 4.25: Heterogeneous treatment effect across previous activity level
Note: Shares of savers switching from the dominated to the dominating fund, across quintiles of years since the most recent investment change. The average years since the most recent investment change for each quintile is denoted below the bars. The average share who switched is depicted in the two leftmost bars.

Year of birth

Heterogeneous treatment effects for the outcome variable switching funds over year of birth quintiles are presented in Figure 4.26. The tendency that the oldest quintile responds relatively strongly is present also for the other outcome variables (Y^l, Y^{mh}, Y^{any}). However, how the youngest group responds relative to the other quintiles varies depending on the outcome variable.

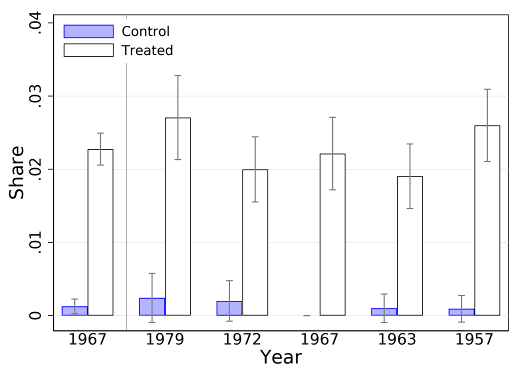

Figure 4.26: Heterogeneous treatment effect across year of birth
Note: Shares of people switching funds, by year of birth quintile, separate for the control group and all treatment groups combined. The mean year of birth in each quintile is shown below the corresponding bars. The left most bars show the averages for all years of birth.

Expected future reward from switching

Figure 4.27 displays heterogeneous treatment effects across quantiles of different specifications of the expected future reward from switching. We conclude that there are small differences between treatment groups where the expected future reward is explicitly stated versus those where only the fee differences between the dominated and dominating funds are presented. In general, there appear to be small differences in responses across expected (discounted) future reward quintiles. This is evident also for the other outcome variables (Y^l, Y^{mh}, Y^{share}). When we normalize the discounted future reward by current labor income, we see that the higher quintiles respond less.

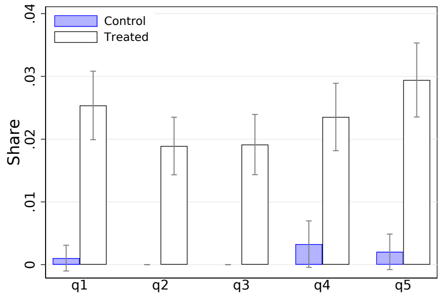

(a) Quintiles of future reward

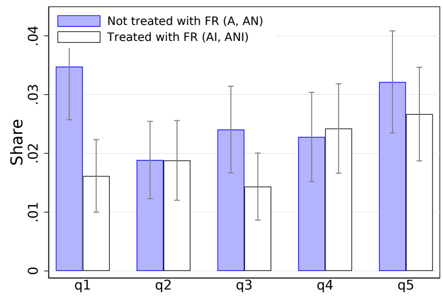

(b) Quintiles of future reward

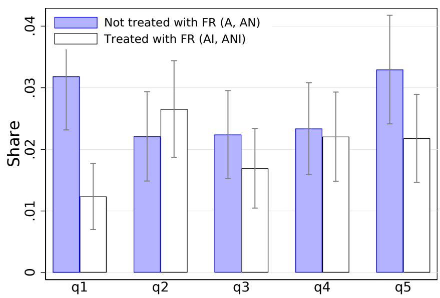

(c) Quintiles of discounted future reward

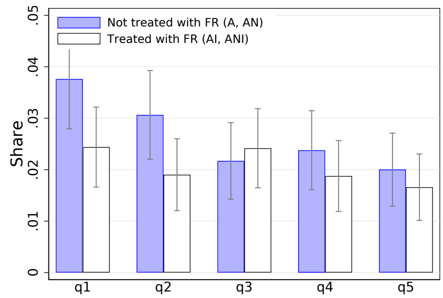

(d) Quintiles of discounted future reward over labor income

Figure 4.27: Heterogeneous treatment effects across future rewards
Note: Shares of savers switching from the dominated to the dominating fund, across (a) quantiles of the expected future reward from switching, all treatment groups vs the control (b) quantiles of the expected future reward from switching, for groups treated with information about the expected reward from switching (AI_0, ANI) vs groups that were not treated with this information (A, AN) (c) quantiles of the discounted expected future reward from switching, for groups treated with information about the expected reward from switching (AI_0, ANI) vs groups that were not treated with this information (A, AN) (d) quantiles of the discounted expected future reward from switching divided by labor income, for groups treated with information about the expected reward from switching (AI_a, ANI) vs groups that were not treated with this information (A, AN). The lines indicate the 95 percent confidence intervals.

Expected total reward from switching

Figure 4.28 presents heterogeneous treatment effects for different quantiles of the total expected discounted future reward from switching, and in the second panel, this measure is normalized by labor income. The total reward amount includes both the expected discounted future reward and the immediate search reward. The figure displays the shares of savers switching funds in treatment group AI, where immediate search rewards are offered.

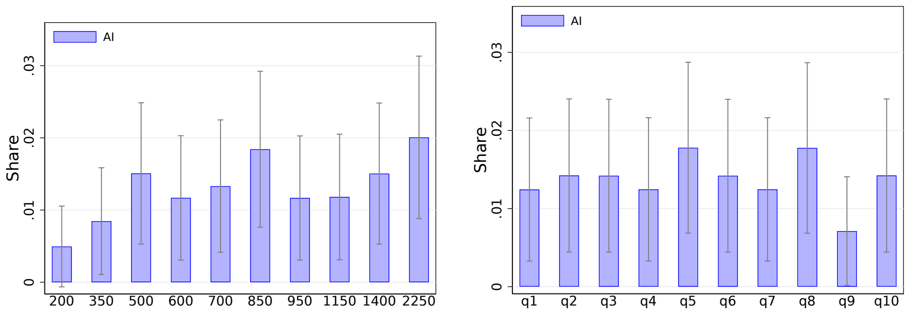

(a) Deciles of total discounted rewards **(b)** Deciles of total discounted rewards over labor income

Figure 4.28: Heterogeneous treatment effects across total discounted rewards
Note: Shares of savers switching from the dominated to the dominating fund, across (a) deciles of the total expected discounted future reward from switching, treatment AI (the mean within each decile is displayed below in SEK, 1 SEK is approximately 0.11 USD) (b) deciles of the total expected discounted future reward from switching over labor income, treatment AI. The vertical lines indicate the 95 percent confidence intervals.

4.E.2 Additional analyses and results

Procrastination and task completeness

We observe the dates when people complete the search task, the reading confirmation, and the verification of the dominating fund name, and *switch funds*. As seen in Figure 4.29 we find no direct signs of procrastination.

Duration of treatment effects

A distribution over the dates of fund switches from the dominated to the dominating fund, following the treatment date is displayed in Figure 4.30. The distributions look similar for the other outcome variables (Y^l, Y^{mh}, Y^{any}). In Figure 4.31, we observe the share of savers that log into their Premium Pension accounts at the Pensions Agency's website, for the treated versus the control group.

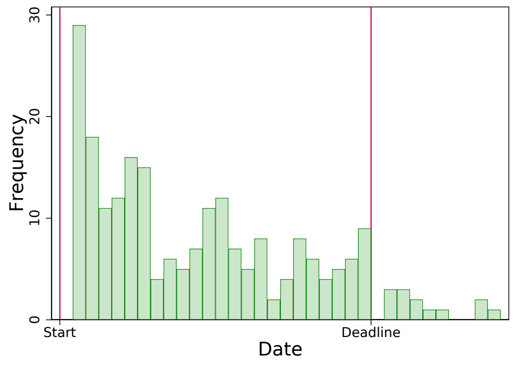

Figure 4.29: Time of task completeness *Note*: Distribution of the time when the search task, the reading confirmation, or the verification of the dominating fund name, was performed, for the people who completed these tasks. The vertical line indicates the deadline to receive the immediate search reward, the reading confirmation compensation, and to perform the name verification. The time from the start to the deadline was 22 days.

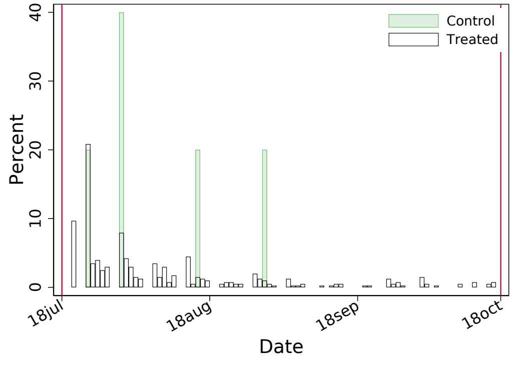

Figure 4.30: Time of fund switches
Note: Distribution of the time of fund switches, following the treatment, shown separately for the control group and all treatment groups combined. The vertical red line indicates the treatment date. Fund changes made during weekends are registered on the following weekday.

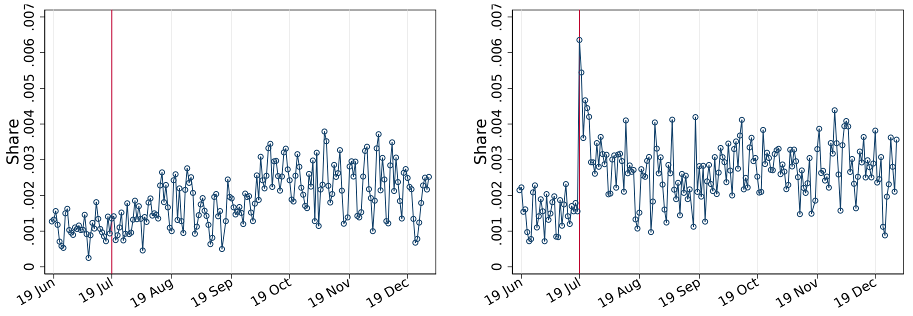

(a) Control group **(b)** All treatment groups

Figure 4.31: Logins at the Premium Pension website

Note: Shares of savers that each day log into their accounts at the Pensions Agency's website, over time. The data is retrieved from the Swedish Pensions Agency in January 2019.

Sammanfattning

Den här avhandlingen består av fyra fristående kapitel. En röd tråd genom avhandlingen är ett fokus på hur olika människor påverkas på skilda sätt av policy. Jag undersöker hur en analys som tar hänsyn till olika hushålls förutsättningar kan påverka slutsatserna i policyutvärderingar samt påverka vår förståelse kring hur policyförändringar påverkar ekonomin i sin helhet. Mer exakt omfattar det jag studerar: ett avskaffande av ränteavdraget på bolån, penningpolitik, mer strikta utlåningsregler för bolån, samt fondval i premiepensionssystemet.

En annan gemensam nämnare i de tre första kapitlen är ett fokus på policy som berör bolåne- och husmarknaden. Skatter och regleringar av bolånemarknaden leder i många länder ofta till intensiva politiska diskussioner. För många hushåll hör beslut kring bostadsköp och bolån till de största och viktigaste finansiella besluten i livet. Dessutom påverkar policy och regleringar i bolånemarknaden ofta olika hushåll på skilda sätt. En utvärdering av en sådan policy kräver därför en analys som kan ta höjd för hur olika människor påverkas. Nedan följer en något mindre teknisk sammanfattning av de fyra uppsatserna.

I kapitel 1, **Kostsamma reformer av dåliga subventioner — fallet med ränteavdraget** (Costly reversals of bad policies: the case of the mortgage interest deduction), tillsammans med Markus Karlman och Kasper Kragh-Sørensen, studerar vi hur hushåll i USA påverkas om man tar bort ränteavdraget för bolån och huruvida detta är en bra idé.

Ränteavdraget är en skattesubvention som har fått en hel del uppmärksamhet i de politiska diskussionerna USA. Subventionen gör det möjligt för husägare att dra av räntebetalningar på hypotekslån från sina skattepliktiga inkomster. Eftersom avdragsrätten för hypotekslån kan minska husägarnas skatter, minskar det i praktiken kostnaden för bolån och därmed kostnaden för att äga en bostad. Sålunda påverkas många hushåll av ränteavdraget, inte bara i sina beslut att äga eller hyra en bostad, men också när det gäller hur stort hus man väljer att köpa. Subventionen kritiseras emellertid ofta för att främst främja höginkomsttagare på andra skattebetalares bekostnad. Ungefär hälften av avdragen går till hushåll i de övre 20 procenten av inkomstfördelningen, medan hushållen i de lägsta 20 procenten knappt gör några ränteavdrag.

För att skapa en bättre förståelse för vem som skulle dra nytta av och vem som skulle förlora på att avskaffa avdragsrätten för räntebetalningar på hypotekslån, utför vi experiment i en modell som är utformad för att representera det amerikanska samhället. Vi börjar med att analysera de långsiktiga välfärdseffekterna, dvs vi jämför om hushållen skulle föredra att födas in i ett samhälle med eller utan avdragsrätt för hypotekslån. Våra resultat visar att en stor majoritet av hushållen skulle föredra ett samhälle utan ränteavdrag. I ett samhälle utan skattesubventionen väljer hushåll med högre inkomster att bo i mindre egenägda bostäder. Detta leder till lägre priser för ägda och hyrda bostäder, vilket är speciellt gynnsamt för hushåll med låga inkomster. Vidare, när regeringen inte längre subventionerar räntebetalningar på hypotekslån kan andra skatter sänkas. Medan enbart vissa hushåll drar nytta av avdragsrätten för hypotekslån gynnas alla hushåll av en lägre inkomstskatt.

Givet de stora välfärdsvinsterna av att ta bort hypotekslånesubventionen på lång sikt, fortsätter vi med att undersöka hur nuvarande hushåll skulle påverkas av ett borttagande. Effekterna av ett avlägsnande är väldigt annorlunda för dessa hushåll. I dag har många hushåll tagit långsiktiga bostads-och bolånebeslut baserat på antagandet att de kan göra ränteavdrag. När subventionen oväntat tas bort faller bostadspriserna kraftigt, vilket drabbar de existerande husägarna avsevärt. Vidare inser många hushåll att de har för stora hus och bolån, när de inte längre kan dra av sina räntebetalningar. De som hyr, å andra sidan, vinner på reformen då de drar nytta av fallet i bostadspriserna.

Våra resultat visar att hushållen i genomsnitt får det sämre om ränteavdraget omedelbart tas bort i sin helhet och en majoritet av hushållen är negativt inställda till en sådan reform. 70 procent av hushållen i USA äger sina hem och de positiva effekter som hyresgästerna får överstiger inte de negativa effekterna för husägarna. Vi visar också att dessa resultat även håller om avskaffandet av ränteavdraget tas bort gradvis eller om det tillkännages i förväg. Våra resultat tyder på att ännu färre hushåll är positiva till ett avskaffande under dessa alternativa implementeringar. Trots att ett mer gradvist borttagande mildrar förlusterna för dem som drabbas värst av reformen, minskar det också vinsterna. Därmed visar våra resultat att kostnaderna för att reformera en dålig politik kan vara avsevärda – även i en sådan utsträckning att det kanske inte är värt det.

I kapitel 2, **Penningpolitik och bolånemarknaden** (Monetary policy and the mortgage market), utforskar jag vikten av förändringar av bolåneräntor och bostadspriser för penningpolitikens effekt. Dessutom undersöker jag om effekterna av en centralbanks räntesättning påverkas av om hushåll har bolån med bunden ränta jämfört med rörlig ränta.

En viktig fråga inom nationalekonomin är hur en centralbank påverkar ekonomin genom att ändra styrräntan. I traditionella modeller som används för att studera penningpolitik leder en sänkning av styrräntan till att hushållen konsumerar mer idag och sparar minde, då en lägre ränta i praktiken sänker priset på konsumtion idag jämfört med i framtiden. Det finns emellertid andra sätt som en lägre ränta påverkar hushåll. I det här kapitlet studerar jag hur hushållens efterfrågan påverkas av en räntesänkning då denna även påverkar bolåneräntor och bostadspriser.

När en centralbank sänker räntan påverkas många hushåll direkt genom lägre räntebetalningar på sina bolån. Vissa hushåll kommer till följd av de lägre utgifterna öka sin konsumtion av andra varor och tjänster, vilket innebär en ökad efterfrågan i ekonomin. Hur mycket hushållens efterfrågan ökar påverkas i sin tur av hur mycket deras räntebetalningar minskar, samt hur finansiellt begränsade hushållen är till att börja med. Om sedan huspriserna också ändras till följd av den lägre styrräntan, kommer bostadsägares förmögenheter att påverkas, vilket också har verkan på deras konsumtionsbeslut.

Jag börjar min analys med att använda en modell av den amerikanska ekonomin, där det vanligaste bolånekontraktet är ett 30-årigt bolån med bunden ränta, vilket innebär fasta betalningar varje år. Med den här typen av kontrakt påverkas bara hushåll som tar upp nya bolån av en sänkning av bolåneräntan. Ett annat bolånekontrakt som är vanligt i många andra länder är ett kontrakt med rörlig ränta. När bolån har rörlig ränta påverkar en räntesänkning även räntebetalningarna för befintliga bolån. Jag fortsätter min analys genom att studera hur efterfrågan i ekonomin påverkas av en sänkt styrränta när bolånekontrakten har rörlig ränta.

Mina resultat visar att i en ekonomi där bolån med bunden ränta används, leder förändringar av bolåneräntor samt huspriser till följd av en sänkt styrränta till en ökad efterfrågan, jämfört med om dessa priser hålls konstanta. En lägre styrränta leder till lägre bolåneräntor samt högre bostadspriser, vilket i sin tur leder till en ökad konsumtion. Hushåll som omförhandlar villkoren på sina bolån ökar sin konsumtion mest. De som omförhandlar sina bolån tenderar att vara i behov av ökade likvida tillgångar, och väljer därför att omförhandla sitt bolån för att dels ta del av den lägre bolåneräntan, men också för att ta ut ett större bolån. Eftersom huspriserna dessutom är högre, kan dessa hushåll ta ut betydligt större bolån, vilket gör att de kan öka sin konsumtion ytterligare.

När istället bolån med rörlig ränta används, visar mina resultat att konsumtionen ökar med mer än sex gånger så mycket till följd av en sänkning av styrräntan, jämfört med om bolånen har bunden ränta. Dels påverkas alla hushåll med bolån av räntesänkningen då räntan på lånen är rörlig; dessutom påverkas den kortare räntan, som används för bolån med rörlig ränta, mer i närtid än den längre räntan som är kopplad till bolån med bunden ränta, då centralbanken temporärt sänker styrräntan. Detta leder till att hushållens bolånebetalningar minskar mer på kort sikt. Därutöver ökar huspriserna betydligt mer i en ekonomi där bolånen har rörlig ränta. Återigen spelar hushåll som väljer att ta ut större bolån till följd av räntesänkningen en stor roll för den ökade efterfrågan. Då huspriserna stiger mer när bolånen har rörlig ränta, kan de som omförhandlar sina bolån ta ut ännu större nya lån, vilket leder till att de kan öka sin konsumtion i en ännu högre utsträckning.

Sammanfattningsvis tyder mina resultat på att hur effektivt en centralbank kan stimulera efterfrågan i ekonomin beror på vilken typ av bolånekontrakt som används, samt tillgängligheten att omförhandla bolånekontrakt. Min slutsats är att det är av stor vikt i penningpolitisk analys att förstå hur bolånemarknaden fungerar och vilka kontrakt som används.

I kapitel 3, **Utlåningsregler för bolån: implikationer för fluktuationer i konsumtion** (Mortgage lending standards: implications for consumption dynamics), även detta samförfattat med Markus Karlman och Kasper Kragh-Sørensen, studerar vi huruvida mer strikta regler för bolån kan minska fallet i konsumtionen under ekonomiska nedgångar. Mer specifikt studerar vi i vilken utsträckning bolåneregler påverkar i vilken omfattning hushåll ändrar sin konsumtion, när de upplever en tillfällig minskning av sina tillgångar.

Myndigheter i många länder har infört striktare krav för bolån under senare år. Denna utveckling är delvis motiverad av erfarenheterna från den stora recessionen, där områden med en högre tillväxt i skuldsättningen via bolån innan krisen upplevde en kraftigare minskning av konsumtionen när krisen slog till. Med de nya bolånekraven hoppas man att framtida nedgångar blir mindre allvarliga. Det är emellertid inte uppenbart att de striktare utlåningskraven är framgångsrika när det gäller att stabilisera ekonomin. Ett sätt på vilket hushållen kan undvika en tillfällig minskning av konsumtionen är just genom att öka sin skuldsättning. Genom att begränsa möjligheterna att låna har hushållen färre möjligheter att mildra konsekvenserna av en minskning av sina finansiella resurser. Därmed kan konsumtionsresponsen till och med vara starkare när striktare regleringar är på plats.

I den här artikeln använder vi en modell för att utföra experiment där belåningsgradskravet (med andra ord kontantinsatskravet) och skuldkvotskravet görs mer strikta. Belåningsgradskravet specificerar det maximala bolånet ett hushåll kan erhålla, som en andel av bostadens värde. Skuldkvotskravet begränsar storleken på bolånet i förhållande till inkomsten. I våra experiment studerar vi först en permanent förskjutning av belåningsgradskravet från det nuvarande värdet på 0,90 till 0,70, eller skuldkvotskravet från dess nuvarande värde på 0,28 till 0,18 (det nuvarande värdet specificerar att inte mer än 28 procent av den årliga inkomsten får läggas på bostadsrelaterade kostnader, dessa kostnader inkluderar räntebetalningar och amortering av bolån). Vi utforskar sedan en temporär implementering av de striktare kraven, under ett år som föregår en ekonomisk nedgång.

Vårt första resultat är att permanent striktare bolåneregleringar enbart marginellt påverkar hur mycket hushållen minskar sin konsumtion, vid en ekonomisk nedgång. De striktare kraven påverkar emellertid hushållen på flera viktiga sätt. Färre hushåll äger sina bostäder, de har lägre skuldsättning och sparar i genomsnitt aningen mer. Av yttersta vikt är dock att dessa beteendeförändringar är sådana att hushållens totala förmåga att hantera ekonomiska nedgångar i princip förblir oförändrad. Det här resultatet håller även för större förändringar av kraven för utlåning.

Vårt andra resultat är att tillfälligt striktare krav för bolån kan framgångsrikt begränsa konsumtionsminskningen under en ekonomisk nedgång. Temporärt striktare krav för bolån förhindrar vissa människor från att köpa hus och leder till att vissa hushåll tar ut mindre bolån. Till följd av detta har hushållen mer disponibla besparingar när den ekonomiska nedgången inträffar än de skulle ha haft utan de striktare regleringarna. Därmed är de bättre förberedda att hantera en minskning av sina tillgångar. Det är emellertid enbart under specifika omständigheter som temporärt striktare bolånekrav leder till att hushållen får det bättre. För det första måste den ekonomiska nedgången vara stor. För det andra behöver en beslutsfattare ha en informationsfördel genom att denne kan förutse nedgången, medan hushållen inte kan göra det.

I det fjärde kapitlet, **Tröghet i pensionsinvesteringar: evidens från ett informationsexperiment** (Inertia of dominated pension investments: evidence from an information intervention), samförfattat med Louise Lorentzon, undersöker vi potentiella anledningar till varför pensionssparare misslyckas med att välja fonder med lägst avgifter, bland fonder med samma investeringsstrategi. Därtill studerar vi om information kring billigare, jämförbara fonder kan förbättra pensionssparares investeringsbeslut.

Trots att det finns ett stort antal fonder för pensionssparare att välja bland, kännetecknas fondmarknaden av bristfällig konkurrens då jämförbara fonder kan ha väldigt olika avgifter. Till och med bland fonder som har exakt samma investeringsstrategi är det stor spridning på avgifterna. Att olika människor betalar olika avgifter för samma placering bidrar till den spridning i avkastning på sparande som vi ser i data. Nya studier visar dessutom att skillnader i avkastning på sparande är en viktig orsak till ojämlikhet i tillgångar.

I det här kapitlet utforskar vi tre potentiella anledningar till varför personer fortsätter att spara i fonder som är dominerade, d.v.s., där det finns en annan fond som har en lägre avgift och som har samma investeringsstrategi. För det första kanske vissa sparare inte är medvetna om att det finns ett billigare jämförbart alternativ. För det andra kanske vissa sparare inte orkar leta efter den billigaste fonden. Att sparare anser att det är för kostsamt att försöka hitta ett billigare alternativ är en förklaring som ofta ges i litteraturen som studerar anledningar till dominerade val. För det tredje kanske sparare missbedömer hur stora vinsterna skulle vara om de byter till en billigare fond, genom att underskatta effekten av ränta på ränta.

För att studera anledningar till dominerade fondval, och för att undersöka om information kan förbättra sparares investeringsval, skickar vi informationsbrev till personer som har valt att spara delar av sin premiepension i någon av två dominerade indexfonder. Våra resultat visar att information som ökar medvetenheten kring att en billigare, jämförbar fond är tillgänglig leder till att flera sparare byter till den billigare fonden. Information där också namnet på den billigare fonden nämns, leder till att en ännu större andel av spararna förbättrar sina investeringar. Information som förtydligar hur mycket mer pengar som spararen kan förvänta sig att ha på sitt premiepensionskonto vid sin pensionering, om denne byter till den billigare fonden, leder emellertid inte till att fler byter fond.

Även om många pensionssparare förbättrar sina investeringsbeslut till följd av informationsbreven fortsätter en majoritet av personerna i studien att spara i de dominerade fonderna. Vi drar slutsatsen att information om jämförbara fonder är användbar för många sparare, och detta kan vara viktigt att ha i åtanke då pensionssystem utformas. Varför så många sparare däremot inte förbättrar sina investeringsbeslut fortsätter att förbrylla, framförallt när sparare får information om namnet på en billigare, jämförbar fond, vilket eliminerar kostnaderna kring att hitta ett bättre alternativ.

IIES monograph series

1. Michaely, Michael, *The Theory of Commercial Policy: Trade and Protection*, 1973
2. Söderström, Hans Tson, *Studies in the Microdynamics of Production and Productivity Change*, 1974
3. Hamilton, Carl B., *Project Analysis in the Rural Sector with Special Reference to the Evaluation of Labour Cost*, 1974
4. Nyberg, Lars and Staffan Viotti, *A Control Systems Approach to Macroeconomic Theory and Policy in the Open Economy*, 1975
5. Myhrman, Johan, *Monetary Policy in Open Economies*, 1975
6. Krauss, Melvyn, *International Trade and Economic Welfare,* 1975
7. Wihlborg, Clas, *Capital Market Integration and Monetary Policy under Different Exchange Rate Regimes,* 1976
8. Svensson, Lars E.O., *On Competitive Markets and Intertemporal Resources Allocation*, 1976
9. Yeats, Alexander J., *Trade Barriers Facing Developing Countries*, 1978
10. Calmfors, Lars, *Prices, Wages and Employment in the Open Economy*, 1978
11. Kornai, Jànos, *Economics of Shortage, Vols I and II*, 1979
12. Flam, Harry, *Growth, Allocation and Trade in Sweden. An Empirical Application of the Heckscher-Ohlin Theory*, 1981
13. Persson, Torsten, *Studies of Alternative Exchange Rate Systems. An Intertemporal General Equilibrium Approach*, 1982
14. Erzan, Refik, *Turkey's Comparative Advantage, Production and Trade Patterns in Manufactures. An Application of the Factor Proportions Hypothesis with Some Qualifications*, 1983
15. Horn af Rantzien, Henrik, *Imperfect Competition in Models of Wage Formation and International Trade*, 1983

16. Nandakumar, Parameswar, *Macroeconomic Effects of Supply Side Policies and Disturbances in Open Economies*, 1985
17. Sellin, Peter, *Asset Pricing and Portfolio Choice with International Investment Barriers*, 1990
18. Werner, Ingrid, *International Capital Markets: Controls, Taxes and Resources Allocation*, 1990
19. Svedberg, Peter, *Poverty and Undernutrition in Sub-Saharan Africa: Theory, Evidence, Policy*, 1991
20. Nordström, Håkan, *Studies in Trade Policy and Economic Growth*, 1992
21. Hassler, John, Petter Lundvik, Torsten Persson, and Paul Söderlind, *The Swedish Business Cycle: Stylized facts over 130 years,* 1992
22. Lundvik, Petter, *Business Cycles and Growth*, 1992
23. Söderlind, Paul, *Essays in Exchange Rates, Business Cycles and Growth*, 1993
24. Hassler, John A.A., *Effects of Variations in Risk on Demand and Measures of Business Cycle Comovements*, 1994
25. Daltung, Sonja, *Risk, Efficiency, and Regulation of Banks*, 1994
26. Lindberg, Hans, *Exchange Rates: Target Zones, Interventions and Regime Collapses*, 1994
27. Stennek, Johan, *Essays on Information-Processing and Competition*, 1994
28. Jonsson, Gunnar, *Institutions and Incentives in Monetary and Fiscal Policy*, 1995
29. Dahlquist, Magnus, *Essays on the Term Structure of Interest Rates and Monetary Policy*, 1995
30. Svensson, Jakob, *Political Economy and Macroeconomics: On Foreign Aid and Development,* 1996
31. Blix, Mårten, *Rational Expectations and Regime Shifts in Macroeconometrics*, 1997
32. Lagerlöf, Nils-Petter, *Intergenerational Transfers and Altruism*, 1997
33. Klein, Paul, *Papers on the Macroeconomics of Fiscal Policy*, 1997
34. Jonsson, Magnus, *Studies in Business Cycles*, 1997
35. Persson, Lars, *Asset Ownership in Imperfectly Competitive Markets*, 1998
36. Persson, Joakim, *Essays on Economic Growth*, 1998

37. Domeij, David, *Essays on Optimal Taxation and Indeterminacy*, 1998
38. Flodén, Martin, *Essays on Dynamic Macroeconomics*, 1999
39. Tangerås, Thomas, *Essays in Economics and Politics: Regulation, Elections and International Conflict*, 2000
40. Lidbom, Per Pettersson, *Elections, Party Politics and Economic Policy*, 2000
41. Vestin, David, *Essays on Monetary Policy*, 2001
42. Olofsgård, Anders, *Essays on Interregional and International Political Economics*, 2001
43. Johansson, Åsa, *Essays on Macroeconomic Fluctuations and Nominal Wage Rigidity*, 2002
44. Groth, Charlotta, *Topics on Monetary Policy*, 2002
45. Gancia, Gino A., *Essays on Growth, Trade and Inequality*, 2003
46. Harstad, Bård, *Organizing Cooperation: Bargaining, Voting and Control*, 2003
47. Kohlscheen, Emanuel, *Essays on Debts and Constitutions*, 2004
48. Olovsson, Conny, *Essays on Dynamic Macroeconomics*, 2004
49. Stavlöt, Ulrika, *Essays on Culture and Trade*, 2005
50. Herzing, Mathias, *Essays on Uncertainty and Escape in Trade Agreements*, 2005
51. Bonfiglioli, Alessandra, *Essays on Financial Markets and Macroeconomics*, 2005
52. Pienaar, Natalie, *Economic Applications of Product Quality Regulation in WTO Trade Agreements*, 2005
53. Song, Zheng, *Essays on Dynamic Political Economy*, 2005
54. Eisensee, Thomas, *Essays on Public Finance: Retirement Behavior and Disaster Relief*, 2005
55. Favara, Giovanni, *Credit and Finance in the Macroeconomy*, 2006
56. Björkman, Martina, *Essays on Empirical Development Economics: Education, Health and Gender*, 2006
57. Larsson, Anna, *Real Effects of Monetary Regimes*, 2007
58. Prado, Jr., Jose Mauricio, *Essays on Public Macroeconomic Policy*, 2007
59. Tonin, Mirco, *Essays on Labor Market Structures and Policies*, 2007
60. Queijo von Heideken, Virginia, *Essays on Monetary Policy and Asset*

Markets, 2007

61. Finocchiaro, Daria, *Essays on Macroeconomics*, 2007
62. Waisman, Gisela, *Essays on Discrimination and Corruption*, 2008
63. Holte, Martin Bech, *Essays on Incentives and Leadership*, 2008
64. Damsgaard, Erika Färnstrand, *Essays on Technological Choice and Spillovers*, 2008
65. Fredriksson, Anders, *Bureaucracy, Informality and Taxation: Essays in Development Economics and Public Finance*, 2009
66. Folke, Olle, *Parties, Power and Patronage*, 2010
67. Drott, David Yanagizawa, *Information, Markets and Conflict*, 2010
68. Meyersson, Erik, *Religion, Politics, and Development*, 2010
69. Klingelhofer, Jan, *Models of Electoral Competition*, 2010
70. Perrotta, Maria Carmela, *Aid, Education and Development*, 2010
71. Caldara, Dario, *Essays on Empirical Macroeconomic*s, 2011
72. Mueller, Andreas, *Business Cycles, Unemployment and Job Search*, 2011
73. von Below, David, *Essays in Climate and Labour Economics*, 2011
74. Gars, Johan, *Essays on the Microeconomics of Climate Change*, 2012
75. Spiro, Daniel, *Some Aspects of Resource and Behavioral Economics*, 2012
76. Ge, Jinfeng, *Essays on Macroeconomics and Political Economy*, 2012
77. Li, Yinan, *Institutions, Political Cycles and Corruption*, 2012
78. Håkanson, Christina, *Changes in Workplaces and Careers*, 2013
79. Qin, Bei, *Essays on Empirical Development and Political Economics*, 2013
80. Jia, Ruixue, *Essays on the Political Economics of China's Development*, 2013
81. Campa, Pamela, *Media Influence on Pollution and Gender Equality*, 2013
82. Seim, David, *Essays on Public, Political and Labor Economics*, 2013
83. Shifa, Abdulaziz B., *Essays on Growth, Political Economy and Development*, 2013
84. Panetti, Ettore, *Essays on the Economics of Banks and Markets*, 2013
85. Schmitt, Alex, *Beyond Pigou: Climate Change Mitigation, Policy Making and Decisions*, 2014

86. Rogall, Thorsten, *The Economics of Genocide and War*, 2015
87. Baltrunaite, Audinga, *Political Economics of Special Interests and Gender*, 2016
88. Harbo Hansen, Niels-Jakob, *Jobs, Unemployment, and Macroeconomic Transmission*, 2016
89. Stryjan, Miri, *Essays on Development Policy and the Political Economy of Conflict*, 2016
90. Karadja, Mounir, *On the Economics and Politics of Mobility*, 2016
91. Kitamura, Shuhei, *Land, Power and Technology, Essays on Political Economy and Historical Development*, 2016
92. Malmberg, Hannes, *Human Capital in Development Accounting and Other Essays in Economics*, 2017
93. Öberg, Erik, *On Money and Consumption*, 2017
94. Lane, Nathaniel, *States of Development: Essays on the Political Economy of Development in Asia*, 2017
95. Prawitz, Erik, *On the Move: Essays on the Economic and Political Development of Sweden*, 2017
96. Iwanowsky, Mathias, *Essays in Development and Political Economics*, 2018
97. Dehdari, Sirus Håfström, *Radical Right, Identity, and Retaliation*, 2018
98. Harmenberg, Karl, *Essays on Income Risk and Inequality*, 2018
99. Kilström, Matilda, *Households' Responses to Policy in Labor and Credit Markets*, 2018
100. Meriläinen, Jaakko, *Essays in Political Economics*, 2019
101. Sigurdsson, Jósef , *Essays on Labor Supply and Adjustment Frictions*, 2019
102. Cocciolo, Serena, *Participatory Governance and Public Service Provision* , 2019
103. Mitrunen, Matti, *Essays on the Political Economy of Development*, 2019
104. Olsson, Jonna, *Work, wealth, and well-being: Essays in macroeconomics*, 2019
105. Darougheh, Saman, *Search and Mismatch*, 2019
106. van Vlokhoven, Has, *On the Cost of Capital, Profits and the Diffusion of Ideas*, 2020

107. Foltyn, Richard, *Essays in Macroeconomics and Household Finance*, 2020
108. Lerva, Benedetta, *Worms, Farms and Schools: Three Essays on Investment Choices in Uganda*, 2020
109. Kragh-Sørensen, Kasper, *On Housing, Mortgages, and Taxation*, 2020
110. Karlman, Markus, *Essays on housing tax treatment, prices, and macroeconomic implications*, 2020

This thesis consists of four self-contained essays.

Costly reversals of bad policies: the case of the mortgage interest deduction studies how households are affected by removing the mortgage interest deduction, and how the welfare effects depend on the design of the removal policy.

Monetary policy and the mortgage market examines the role of changes in mortgage interest rates and house prices for the effectiveness of monetary policy.

Mortgage lending standards: implications for consumption dynamics studies to what extent stricter mortgage regulations can dampen the fall in consumption during economic downturns.

Inertia of dominated pension investments: evidence from an information intervention investigates why savers fail to choose the pension fund with the lowest fee among funds with the same investment strategy.

Karin Kinnerud
holds a B.Sc. in Business Administration, International Business and a M.Sc. in Quantitative Financial Economics from Oklahoma State University.

ISBN 978-91-7911-206-6
ISSN 0346-6892

Department of Economics

Stockholm University