

Essays on Dynamic Macroeconomics: From Structural Parameter Estimation to the Evaluation of Central Bank Policies

Juan Castellanos Silván

January, 2025

Abstract

The thesis is composed by three independent essays on Dynamic Macroeconomics. The first chapter is related to the estimation of structural parameters within a DSGE model. A Monte Carlo study is implemented to examine the small sample performance of IRF matching and Indirect Inference estimators that target impulse responses (IRFs) that have been estimated with Local Projections (LP) or Vector Autoregression (VAR). The analysis considers various identification schemes for the shocks and several variants of LP and VAR estimators. Results show that the lower bias from LP responses is a big advantage when it comes to IRF matching, while the lower variance from VAR is desirable for Indirect Inference applications as it is robust to the higher bias of VAR-IRFs. Overall I recommend the use of Indirect Inference over IRF matching when estimating DSGE models as the former is robust to potential misspecification coming from identification assumptions, small sample or incorrect lag selection.

In the second chapter, joint with Andrew Hannon (ECB) and Gonzalo Paz-Pardo (ECB), we build a model of the aggregate housing and rental markets in which house prices and rents are determined endogenously. Households can choose their housing tenure status (renters, homeowners, or landlords) and the size of their homes depending on their age, income and wealth. We use our model to study the impact of changes in credit conditions on house prices, rents and household welfare. We analyze the introduction in Ireland in 2015 of policies that limited loan-to-value (LTV) and loan-to-income (LTI) ratios of newly originated mortgages and find that, consistently with empirical evidence, they mitigate house price growth but increase rents. Homeownership rates drop, and young and middle-income households are negatively affected by the reform. An unexpected permanent rise in real interest rates has similar effects – by making mortgages more expensive and alternative investments more attractive for landlords, it increases rents relative to house prices.

The third chapter, joint with Stephen Millard (NIESR) and Alexandra Varadi (BoE), offers a structural approach to the following questions: *how does the strength of mone-*

tary policy depend on the mortgage interest fixation period? And how it is affected by credit conditions? In many countries the most common mortgage contract neither has a fixed nor a fully adjustable rate. The typical interest fixation period varies between two to ten years. We embedded such a contractual arrangement into a DSGE with long term mortgage debt and loan-to-value (LTV) and payment-to-income (PTI) constraints to study the effects of monetary policy as well as its interaction with those borrower-based macro-prudential limits. After calibrating the model to the United Kingdom, we find that: (i) the interest fixation period and the tightness of credit conditions do not matter if monetary policy shocks are transitory, (ii) looser credit limits and shorter fixation periods amplify the redistributive effects of inflation target shocks that increase nominal rates persistently, and (iii) LTV limits act as a backstop to the high sensitivity of PTI limits to monetary policy, especially when the interest fixation period is short.

Acknowledgements

This thesis, even though it is composed by three chapters, closes a bigger one in my academic and personal life. During this period I have greatly benefited from the interaction with many fantastic human beings.

First, I would like to start by thanking my wonderful advisors Russell Cooper and Ramon Marimon for their continuous guidance, support, encouragement and patience as my Ph.D. journey has been everything but standard. After a challenging first year, in which Adrien, Anna, Leo and Marcin were the best problem set team one can ask for, I visited the University of Pennsylvania (UPenn) with a La Caixa Foundation Fellowship whose financial support is greatly acknowledged. I am grateful to UPenn's Department of Economics for their hospitality, and in particular, I want to express my gratitude to all the Professors that allowed me to take part in their advanced courses. Moreover, I am forever indebted to Ramon for not only organizing this visit but also vouching for me in front of the ECO Department and the Spanish Grant Authority at the European University Institute (EUI). I would also like to thank José Antonio, Michela, Fatma, and Sarah for helping me navigate all the bureaucratic issues in such daunting process.

After ten intense but productive months in Philadelphia, I returned to Florence and started to work closely with Russell as teaching and research assistant. I felt very privileged to spend so many hours in his office discussing the in and outs of dynamic macroeconomics. That period enormously shaped my identity as a researcher, and in fact, the first chapter of this thesis was motivated by a bonus question in one of the problem sets in his *Statistics and Econometrics III* course! Russell was initially a co-author in the project but he let me fly solo once I was mature enough to single handed the project. What else can you ask from a supervisor?

In my fourth year I flew the nest, direction to Frankfurt, the seat of the European Central Bank. There I was privileged to interact with great economists such as Ivan Jaccard, Caterina Mendicino, Gonzalo Paz-Pardo, Marcel Peruffo, Valerio Scalone, Sebastian Schmidt, Jirka Slacalek or Alejandro Van der Ghote, from whom I learnt a lot.

To Gonzalo I am particularly grateful as he took me into a project with Andrew Hannon that has become the second chapter of this thesis. These past three years have been a fantastic learning experience on how to shape a project into a polished paper. If Russell and Ramon are my academic fathers, then Gonzalo must be my academic big brother. In fact, the ECB community felt like family away from home. I enjoyed a lot my daily interactions with the fascinating Trainee community, from DGR's 7.20 office to the IPA extended family.

After one last dance in Florence, which I cherish for all the basketball conversations and games played with the "Tres Amigos" and that help me recharge the batteries after long working days, I went to my last Ph.D. visit. This time at the Bank of England. In the three months I was there, we – Alexandra Varadi, Stephen Millard and I – started to work in what turned out to be the third and last chapter of my thesis. I am thankful for the time they spent working on that project and in particular to Alex for being my cicerone at the Bank. She is a big part of why I returned and now work at Bank of England. Moreover, I want to thank all "the Arbitrageurs" and my cohort of PhD Interns for the pints had while discussing economics. As a Professor of mine used to say: "If you get stuck with an idea/project, go get a beer". It works like a charm as it enhances creativity.

I would also like to express my gratitude to the external board members, Martin Eichenbaum and Gianluca Violante, for their effort and thoughtful comments. Moreover, my work has also benefited from discussions with other faculty members at the EUI like Jesus Bueren, Andrea Gazzani, Edouard Challe, and Barbara Rossi, as well as from multiple discussions with conference participants across the world. Thank you very much for your time.

I cannot finish these lines without thanking my family and friends. My Madrid oasis for always being there to celebrate life but also to host me whenever I needed a place to stay during my various travels. My high school friends for all the times I told you I wasn't an accountant or a financial adviser. My American family, Nick, Nancy and Tom, for treating me like your own since we met back in the summer of 2011. My aunt and uncle, Vanesa and Norberto, for all the hikes that helped me clear my mind. My grandmothers, Carmen and Geli, for being my biggest fans. And of course, to my mom and dad, Elvia and Juan Carlos, to whom I dedicate this thesis. Since there are no words to describe how essential you have been in my development, not only as a kid but also as an adult, at least this thesis will put a number to it.

Contents

1	Local Projections vs. VARs for Structural Parameter Estimation	11
1.1	Introduction	12
1.2	The Data Generating Process	15
1.2.1	The Smets and Wouters Model	16
1.3	Estimation Methods	16
1.3.1	Indirect Inference	16
1.3.2	Impulse Response Function Matching	17
1.3.3	The (Auxiliary) Econometric Models	18
1.3.3.1	VAR approaches	18
1.3.3.2	Local projection approaches	19
1.3.3.3	Lag length selection	19
1.3.4	Impulse Response Estimands & Identification	20
1.3.4.1	Observed innovation / observed shock identification	20
1.3.4.2	Recursive identification	21
1.3.4.3	Noisy direct measures of the shocks of interest	22
1.4	Design, Implementation & Evaluation	23
1.4.1	Performance Metrics	25
1.4.1.1	Overall performance	25
1.4.1.2	Parameter-by-parameter performance	26
1.5	Results	27
1.5.1	The Best Case Scenario: Observed Innovations as Benchmark	27
1.5.2	The Good Old-Fashioned Days: Recursive Identification	31
1.5.2.1	Technology shock	31
1.5.2.2	Monetary policy shock	33
1.5.3	Direct Proxies for the Shocks: Measurement Error & Unit Effect Normalization	35

1.5.3.1	Classical measurement error in the innovation	35
1.5.3.2	Correlated measurement error: government spending and its correlation with technology	37
1.5.3.3	Unit normalization: a 1% increase in the policy rate . . .	37
1.6	Conclusion	38
2	The Aggregate and Distributional Implications of Credit Shocks on Housing and Rental Markets	41
2.1	Introduction	42
2.2	The Model Economy	47
2.2.1	Production	48
2.2.2	Households	50
2.2.3	Equilibrium	52
2.2.4	Model intuition: a supply & demand explanation	54
2.3	The Irish macro-prudential reform	55
2.3.1	Empirical evidence	56
2.3.2	Model calibration	58
2.3.2.1	Externally calibrated parameters	59
2.3.2.2	Internally calibrated parameters, targets, and model fit	62
2.3.3	Aggregate and distributional effects of tighter borrowing limits .	64
2.3.3.1	Steady state comparison	65
2.3.3.2	Transition dynamics & welfare	66
2.4	Interest rates, credit standards, and price dynamics	70
2.4.1	A permanent rise in the real interest rate	70
2.4.1.1	An increase in the return on savings versus a rise in the borrowing rate	70
2.4.2	Interaction with credit standards	72
2.4.3	Transition dynamics & welfare	73
2.5	Conclusion	75
3	The Role of Mortgage Interest Fixation Periods for Monetary & Macro-prudential Policies	77
3.1	Introduction	78
3.2	Mortgage Structure: Is the UK Market that Different?	83
3.3	The Model Economy	85
3.3.1	Households	85

3.3.2	Production	88
3.3.3	Monetary authority	90
3.3.4	Key equilibrium conditions	90
3.4	Calibration	92
3.4.1	Externally calibrated parameters	92
3.4.2	Internally calibrated parameters, targets, and model fit	95
3.5	Results	96
3.5.1	Monetary policy pass-through and its effects on consumption	96
3.5.1.1	The length of the fixation period and its impact on consumption	100
3.5.2	Credit limits and the monetary policy transmission	101
3.5.2.1	Alternative LTV and PTI limits: loose vs. tight credit	101
3.5.2.2	The complementarity between LTV and PTI limits	104
3.6	Conclusion	106
A	Appendix to Chapter 1	119
A.1	The Smets-Wouters Model	119
A.2	Data Moments / Targets	121
A.2.1	Observed Innovation	121
A.2.1.1	The bias-variance trade off	121
A.2.1.2	Observed innovation vs. observed shock	122
A.2.2	Recursive Identification	122
A.2.2.1	Technology shock	122
A.2.2.2	Monetary policy shock	123
A.2.3	Direct measures of the shocks of interest	124
A.2.3.1	Uncorrelated external proxies	124
A.2.3.2	The correlated government spending shock	125
A.2.3.3	Unit normalization with uncorrelated external proxies	126
A.3	Hyperparameter Choices	128
A.3.1	Lag Length	128
A.3.2	Sample Size	130
A.3.3	Weighting Matrices	133
B	Appendix to Chapter 2	135
B.1	Additional empirical evidence	135
B.1.1	Irish rental sector	135

B.1.2	Macro-prudential limits, house & rental prices	137
B.1.2.1	Data sources	137
B.1.2.2	Non-parametric evidence	138
B.2	Further model details	139
B.2.1	Solution method	139
B.2.1.1	Household problem	139
B.2.1.2	General equilibrium	140
B.2.1.3	Transition dynamics	141
B.2.2	LTI and LTV implementation in Ireland	142
B.3	Additional model results and experiments	144
B.3.1	Understanding house and rental price responses	144
B.3.2	The macro-prudential reform: interaction between limits	146
C	Appendix to Chapter 3	147
C.1	Further Model Details	147
C.1.1	Fixed & adjustable rate mortgage economies	147
C.1.1.1	FRM economy	148
C.1.1.2	ARM economy	150
C.1.2	The hybrid rate mortgage economy	151
C.1.2.1	The simplest example: a one-period HRM	151
C.1.2.2	A general T-periods HRM	152
C.2	Additional Figures	154
C.2.1	Temporary monetary policy shock	154
C.2.2	Persistent inflation target shock	155
C.2.3	Alternative PTI & LTV calibrations	156
C.2.4	Constraint Switching Effect	157

Chapter 1

Local Projections vs. VARs for Structural Parameter Estimation

Abstract: This paper conducts a Monte Carlo study to examine the small sample performance of IRF matching and Indirect Inference estimators that target impulse responses (IRFs) that have been estimated with Local Projections (LP) or Vector Autoregressions (VAR). The analysis considers various identification schemes for the shocks and several variants of LP and VAR estimators. Results show that the lower bias from LP responses is a big advantage when it comes to IRF matching, while the lower variance from VAR is desirable for Indirect Inference applications as it is robust to the higher bias of VAR-IRFs. Overall, I recommend the use of Indirect Inference over IRF matching when estimating DSGE models as the former is robust to potential misspecification coming from invalid identification assumptions, small sample issues or incorrect lag selection.

1.1. Introduction

The Local Projections (LP) approach to understanding the dynamic effects of exogenous shocks, originated in Jordà (2005), has become a common and alternative tool to the traditional Vector Autoregression (VAR) approach. In light of the theoretical result in Plagborg-Møller and Wolf (2021), i.e. VARs and LPs estimating the same impulse responses in population, one may think that using either VAR or LP should not matter. However, the finite sample properties of these two estimators differ. In particular, when p lags of the data are included in the VAR and as controls in the LP, IRFs approximately agree out to horizon p , but at longer horizons $h > p$ there is a bias-variance trade-off (Li, Plagborg-Møller, and Wolf 2023).

This paper explores the implications of using one of these two econometric models for summarizing key features of the data, such as the dynamic response to exogenous shocks, in order to estimate the structural parameters of a DSGE model. Given their different small sample properties, targeting impulse responses (IRFs) estimated by LP or VAR will lead to different structural parameter estimates. Hence, in practice, using LP- or VAR-IRFs for DSGE estimation may lead to different outcomes and potentially different quantitative predictions from the structural model.

I carry out a Monte Carlo analysis to investigate the consequences of targeting LP vs. VAR estimated responses in a minimum distance estimation. I consider two estimators within this class: impulse response function matching (*IRF matching*) and indirect inference (*Ind. Inf.*), and use the Smets and Wouters' (2007) model as the Data Generating Process (DGP). Further, and as a benchmark, I assume that the econometrician observes the true shock, which guarantees correct identification. Nevertheless, estimated responses will vary depending on the econometric model used for estimation as well as on the sample size and the number of lags. In general, targeting LP responses which have a lower bias than VARs is a great idea if resorting to *IRF matching*. On the other hand, when estimating the structural parameters via *Ind. Inf.*, using the VAR as the auxiliary model outperforms LPs because *Ind. Inf.* is robust to misspecification and VARs have a lower variance.

These results are better understood in conjunction with the choice of p , the lag length. Note that LP responses are independent of the lag length when the shock is observed, while VAR responses become more similar to LP's as the lag length increases. Actually, the reduction of the bias in VAR responses as p increases comes also at the cost of a larger variance (Olea, Plagborg-Møller, Qian, and Wolf 2024). Consequently,

when p is large, there are little differences between targeting LP- or VAR-IRFs. On the other hand, when p is small, the LP approach is significantly better than VARs for *IRF matching* due to its smaller bias, while using a small p VAR as the auxiliary model for *Ind. Inf.* is the superior choice due to its smaller variance.

The sample size used to estimate these responses also has an impact on the structural parameter estimation. The larger the small sample bias, the worse the estimation outcome. However, such deterioration in the performance of the estimation is more important for *IRF matching* than for *Ind. Inf.* to the point that the latter is preferred regardless of the econometric model used to estimate IRFs. Moreover, I also show that using the bias corrected version of LPs and VARs improves the estimation outcome in an *IRF matching* exercise, while it is not so relevant for *Ind. Inf.* applications.

On a second set of Monte Carlo simulations I relax the observed shock assumption and consider a scenario in which the econometrician does not observe the shock at all and has to infer it from recursive assumptions. Here I show that if assumptions are correct, e.g. by assuming that TFP does not affect other endogenous variables at time 0 as it is the case in the Smets and Wouters' model, the results from the observed shock scenario still hold. On the other hand, when these recursive assumptions are incorrect, e.g. if I assume that the policy rate has no contemporaneous impact on real variables, a common assumption for monetary policy shocks that doesn't hold in the Smets and Wouters' model, then *IRF matching* estimates are significantly worse relative to the observed shock identification due to the larger bias in IRFs, while *Ind. Inf.* estimates are surprisingly better than the observed shock case because of the lower variance in IRFs, specially at shorter horizons.

In the last set of Monte Carlo simulations I consider an intermediate scenario in which the econometrician observes a proxy for the shock that is contaminated with measurement error, which can or cannot be correlated with other shocks. In either case, the estimation performance worsens for both (auxiliary) econometric models (LP & VAR) and estimation approaches (*IRF matching* & *Ind. Inf.*). However, an improvement can be attained if applying the unit effect normalization of Stock and Watson (2018), which corrects for the bias in the estimated IRFs, and consequently, improves the structural estimation outcome for both approaches, but specially for *IRF matching*.

Overall, these findings provide a novel perspective on DSGE estimation setups that target estimated impulse responses and shed light on how the bias-variance trade off between LPs and VARs translate to the structural parameters of the economic model. The main lesson is that *Ind. Inf.* is robust to misspecification, which is more common

among VARs, and benefits from more tightly estimated IRFs. The opposite is true for the *IRF matching* approach. Thus, researchers should rely more often on *Ind. Inf.* to estimate their DSGE models as it is robust to potential misspecification coming from invalid identification assumptions, small sample bias or incorrect lag selection.

Related Literature. The Monte Carlo study in this paper is inspired by the seminal work of Smith (1993) on the use of VAR models as the binding function in an indirect inference exercise that estimates the structural parameters of a DSGE model. Unlike Smith (1993), who uses all the coefficients in the VAR, I only select those coefficients that identify the impulse responses to a given shock. Hence, my paper is also related to the literature that relies on IRF matching for DSGE estimation (Rotemberg and Woodford 1997). In fact, throughout the paper, I compare the performance of these two estimators, *Ind. Inf.* and *IRF matching*, when targeting responses to various shocks under different IRF estimation methods and identification strategies. Consequently, my paper belongs to the broader literature that studies the small sample properties of minimum distance, simulation based, partial information estimators. Examples include: (i) Jordà and Kozicki (2011) who propose an estimator in which the economic model restrictions are based on its impulse response representation; (ii) Creel and Kristensen (2011) who propose an Indirect Likelihood Estimator as an alternative to Simulated Method of Moments or Indirect Inference; (iii) Scalone (2018) who advocates for the use of Bayesian Method of Moments for the estimation of non-linear economic models; or (iv) Ruge-Murcia (2007, 2012, 2020) who studies the small sample properties of minimum distance estimators in linear and non-linear environments as well as with linear and non-linear binding functions for the indirect inference applications. Unlike these papers, my Monte Carlo study aims to analyze the small sample properties of the two most common minimum distance estimators used in macroeconomic applications, *IRF matching* and *Ind. Inf.*, under various identification assumptions for the estimated responses that act as targets. Moreover, I consider LPs, in addition to VARs, as the auxiliary econometric model adopted for estimation.

Given my interest in the performance of LP and VARs as the source of (data) moments or as the auxiliary model for indirect inference, my paper is also related to the literature that studies the performance of these two methods in the context of IRF estimation. Plagborg-Møller and Wolf (2021) have proven that these econometric models estimate the same IRFs in population and that LPs can impose the same amount of identification restrictions used in SVARs after appropriately choosing the set of controls. However, the small sample properties of these two estimators differ. In fact, there is a bias-variance

trade-off beyond horizon p as shown in Li, Plagborg-Møller, and Wolf (2023). My paper complements their results as it confirms the bias-variance trade-off under a different DGP, but more importantly, investigates its implications for the purposes of uncovering structural DSGE parameters.

Overview. The rest of the paper is organized as follows. Section 1.2 describes and justifies the choice of the DSGE model used to generate the data. Section 1.3 describes the estimation methodology, the auxiliary models employed to estimate IRFs, and the various identification strategies within the context of the DSGE model used as DGP. Section 1.4 outlines the Monte-Carlo design and discusses the metrics used to evaluate the results, which are then presented in Section 1.5. Finally, Section 1.6 concludes.

1.2. The Data Generating Process

This section describes the model used to generate the data for the Monte Carlo study in which I compare the impulse response function matching (*IRF matching*) and the indirect inference (*Ind. Inf.*) estimation strategies as a way to infer the structural parameters of a DSGE model. Many models could have fulfilled this purpose, nonetheless, I have chosen the Smets and Wouters (2007) model for several reasons. First, it is a well-understood and widely used model in academia as well as in policy circles. Second, the vast majority of existing applications that estimate their model economies by matching impulse responses concern linearized models, see for example Rotemberg and Woodford (1997), Christiano, Eichenbaum, and Evans (2005), Iacoviello (2005), or Jordà and Kozicki (2011). It is true, however, that the theoretical foundations of indirect inference were grounded on the estimation of nonlinear models (Gourieroux, Monfort, and Renault 1993). I acknowledge this limitation associated to the chosen DGP, nonetheless, how to choose between LPs and VARs within these two estimation set-ups is still an open question in linearized, and hence, simpler settings.¹ And third, the model is sufficiently rich to allow us to explore different types of shocks and identification strategies. As discussed in Ramey (2016), monetary, fiscal and technology shocks are the most widely studied in empirical applications and hence responses to these shocks are potentially also being used as data moments/targets for structural estimation. Importantly, the Smets and Wouters model is able to generate reasonable responses to all these three shocks.

¹ Ruge-Murcia (2020) studies the performance of non-linear auxiliary models in non-linear settings, but he only focuses on local projections and indirect inference estimation.

1.2.1. The Smets and Wouters Model

The model is based on Christiano, Eichenbaum, and Evans (2005) who added various frictions to a basic New Keynesian DSGE in order to capture the dynamic response to a monetary policy shock as measured by a structural vector autoregression (SVAR). When price and wage stickiness are paired with adjustment costs for investment, capacity utilization costs, habit formation in consumption, partial indexation of prices and wages as well as autocorrelated disturbance terms, the model is able to generate a rich autocorrelation structure. These elements are key for capturing the joint dynamics of output, consumption, investment, hours worked, wages, inflation and the interest rate in the Euro Area (Smets and Wouters 2003). The 2007 version of the model, which I use in this paper, is a minor modification of the 2003 Smets and Wouters model in order to fit the US data. Given the importance of the Smets and Wouters model in the DSGE literature, I do not describe their economy in this paper. Nonetheless, the log linearized equilibrium conditions are reproduced in Appendix A.1.

1.3. Estimation Methods

1.3.1. Indirect Inference

Any economic model, including the Smets and Wouters model, can be represented as a function, $M(\cdot)$, that for a given vector of parameters Θ maps a sequence of endogenous states $\{y_{t-1}\}$, exogenous variables $\{x_t\}$ and shocks $\{\varepsilon_t\}$, into a sequence of endogenous variables $\{y_t\}$. That is,

$$(1.1) \quad y_t = M(y_{t-1}, x_t, \varepsilon_t; \Theta)$$

for $t = 1, \dots, T$. Therefore, using this mapping it is possible to generate infinite data sequences $\{y_t\}_{t=1}^T$, given an initial value of the endogenous state y_0 and a sequence of the shocks $\{\varepsilon_t\}_{t=1}^T$. Model simulation is in fact the basis for the class of minimum distance estimators that seek to find an ex-ante unknown parameter vector that minimizes the distance between data and simulated moments. The most common among this class are the Simulated Method of Moments (SMM) and the Indirect Inference (*Ind. Inf.*) estimators. The only difference between these two is that SMM uses unconditional moments, while in an *Ind. Inf.* exercise these come from an auxiliary, typically econometric, model. The auxiliary model that summarizes the key features of the data into a tractable vector of parameters is often referred to as the *binding function*.

The *Ind. Inf.* approach was popularized in macroeconomics by Smith (1993) who used a VAR to summarize the key features from the data that he wanted to replicate with his economic model. Formally, the *Ind. Inf.* estimates arise from solving the following minimization problem

$$(1.2) \quad J^{smm} = \min_{\Theta} (\beta - \beta(\Theta))' W (\beta - \beta(\Theta))$$

where β and $\beta(\Theta)$ are the estimated coefficients of an auxiliary (econometric) model and W is a weighting matrix. In this paper, these estimated coefficients are those identifying the dynamic response to an aggregate shock. As shown in Sections 1.3.3 and 1.3.4, there are various approaches to identify and estimate IRFs. Consequently, the main objective of this paper is to study how the choice of these particular binding functions $\beta(\cdot)$ affect the structural parameter estimates $\hat{\Theta}$.

1.3.2. Impulse Response Function Matching

An alternative to *Ind. Inf.*, that has been used extensively in DSGE estimation, is impulse response function matching (*IRF matching*). It is also a minimum distance estimator as it minimizes the distance between data targets (estimated IRFs) and its model counterparts (structural IRFs). As it has been pointed out in Christiano, Eichenbaum, Vigfusson, Kehoe, and Watson (2006), this approach is not fully statistically sound as it compares the small sample distribution of the estimator to the population impulse responses which may have a different asymptotic distribution. In other words, *IRF matching* is minimizing the distance between two statistically different objects, while *Ind. Inf.* approximates the small sample distribution of the targets through simulation and hence has a better statistical foundation. Nonetheless, *IRF matching* provides a natural benchmark to compare the *Ind. Inf.* estimation results because bias-variance trade-offs, small sample biases or incorrect identification strategies associated to the estimated dynamic responses will only affect the data moments/targets. On the other hand, these may cancel each other off when the same model is applied to both data and simulated moments as it is the case in the *Ind. Inf.* approach.

Formally, the *IRF matching* estimates are obtained after solving the following minimization problem

$$(1.3) \quad J^{irf} = \min_{\Theta} (\beta - IRF(\Theta))' W (\beta - IRF(\Theta))$$

where the only difference with respect to (1.2) is on how impulse responses are computed when the candidate vector of parameters Θ is updated in search of a minimum. Notice that in (1.3), the model counterpart, $IRF(\cdot)$, is the structural IRFs and hence they do not require a simulated dataset because they are directly computed from the ABCD representation of the model (Fernández-Villaverde, Rubio-Ramírez, Sargent, and Watson 2007). To be clear, $IRF(\Theta)$ are a function of the structural parameters only, while $\beta(\Theta)$ are also a function of the sample size T , and the lag length p when estimated with a VAR. And in particular, the following is satisfied: $\beta(\Theta) \rightarrow IRF(\Theta)$ as $T \rightarrow \infty$ and $p \rightarrow \infty$ when estimated with a VAR, which allows me to study how *truncation* ($p < \infty$) and *small sample* ($T < \infty$) biases affect the structural parameter estimates (Chari, Kehoe, and McGrattan 2008).

1.3.3. The (Auxiliary) Econometric Models

Assume that I observe data $w_t = (r'_t, \tilde{x}_t, \tilde{y}_t, q'_t)$ where $\tilde{x}_t$ and $\tilde{y}_t$ are scalar time series and r'_t and q'_t are $n_r \times 1$ and $n_q \times 1$ vectors of time series including contemporaneous and lagged controls, respectively. I am interested in the dynamic response of $\tilde{y}_t$ after an impulse in $\tilde{x}_t$ as a way of summarizing some features of the data that I would like to replicate with my structural macroeconomic model. The most common approaches to estimate these impulse responses in the data involve the use of VAR or LP. The choice between these two econometric models is important because, despite estimating the same responses in population (Plagborg-Møller and Wolf 2021), their small sample performance is characterized by a bias-variance trade off (Li, Plagborg-Møller, and Wolf 2023). Hence, I am interested in how these small sample properties may affect the structural parameters when VARs or LPs are used to summarize the data in a minimum distance estimation.

1.3.3.1. VAR approaches

Least Squares VAR. I consider a recursive VAR specification in w_t

$$(1.4) \quad w_t = c + \sum_{\ell=1}^p A_{\ell} w_{t-\ell} + u_t$$

where u_t is the projection residual and $(c, \{A_{\ell}\}_{\ell=1}^p)$ are the projection coefficients. These coefficients are estimated by least-squares and the residual covariance matrix, $\hat{\Sigma}_u = T^{-1} \sum_{t=2}^T \hat{u}_t \hat{u}'_t$, is factorized using a lower triangular Cholesky factor $\hat{B}$, such that $\hat{B} \hat{B}' = \hat{\Sigma}_u$.

Define the lag polynomial $\sum_{\ell=0}^p C_{\ell} L^{\ell} = C(L) \equiv A(L)^{-1}$. Noting that $\tilde{x}_t$ and $\tilde{y}_t$ are the $(n_r + 1)$ -th and the $(n_r + 2)$ -th elements in w_t , I can now define the VAR impulse response function of $\tilde{y}_t$ with respect to an impulse in $\tilde{x}_t$ as $\{\Lambda_h\}_{h \geq 0}$ where

$$(1.5) \quad \Lambda_h \equiv C_{n_r+2, \bullet, h} B_{\bullet, n_r+1}$$

and $B_{\bullet, n_r+1}$ is the $(n_r + 1)$ -th column of B and $C_{n_r+2, \bullet}$ refers to the $(n_r + 2)$ -th row of C_h .

Bias corrected VAR. The impulse responses are estimated as above, but I use the modification proposed by Kilian (1998) that applies the formula in Pope (1990) to analytically correct the bias of the reduced-form coefficients caused by persistent data.

1.3.3.2. Local projection approaches

Least Squares LP. The least-squares local projection estimator β_h is obtained from the OLS regression

$$(1.6) \quad \tilde{y}_{t+h} = \mu_h + \beta_h \tilde{x}_t + \gamma_h' r_t + \sum_{\ell=1}^p \delta_{h,\ell}' w_{t-\ell} + \xi_{h,t}$$

where $\tilde{y}_{t+h}$ is the response variable, $\tilde{x}_t$ is the impulse variable, and r_t are contemporaneous controls, $\{w_{t-\ell}\}_{\ell=1}^p$ controls for p lags of all data series included in the regression, and $\xi_{h,t}$ is the projection residual.

Bias Corrected LP. I use the version proposed by Herbst and Johannsen (2023) where they partially remove the bias associated to high persistence in the data. This bias, although asymptotically negligible relative to the standard deviation, can be sizable in small samples.

1.3.3.3. Lag length selection

A key element to understand the differences in the estimated IRFs when using Local Projections or SVARs is the lag length, p . Recall one of the Plagborg-Møller and Wolf's (2021) results: *Local Projections with p lags as controls and VAR(p) estimators approximately agree at impulse response horizons $h \leq p$* . Consequently, using longer lag lengths given a fixed horizon H will certainly deliver more similar targeted responses across the two econometric models. As a result, the estimated economic parameters should also be

more similar when comparing across VARs and LPs as the source of moments/targets. To test this hypothesis, we will consider estimation setups with various lag lengths for the (auxiliary) econometric models, i.e. I let $p \in \{2, 4, 8, 12\}$. Moreover, this type of exercise is useful to address the fact that VARs with short lag lengths capture poorly DSGEs dynamics, an issue that is not applicable to LPs.

Alternatively, I could have opted for using information criteria such as AIC or BIC, however, these tend to select very short lag lengths which are not consistent with the typical choices in applied work. In fact, Li, Plagborg-Møller, and Wolf (2023) use the following lag length rule, $p = \max\{\hat{p}_{AIC}, 4\}$, which for my DGP will have always resulted in picking $p = 4$.

1.3.4. Impulse Response Estimands & Identification

I follow Li, Plagborg-Møller, and Wolf (2023) in considering three types of structural impulse response estimands to mimic as closely as possible the schemes used in applied macroeconometrics to identify impulse responses in the data. Recall that these responses are simply a way of summarizing the data for our structural estimation, and not the main focus of our analysis.

1.3.4.1. Observed innovation / observed shock identification

I assume that the econometrician observes the endogenous variables $\bar{w}_t$ and the true structural shock ε_t or equivalently its innovation η_t . For the VAR approaches, I order the shock as the first variable in the VAR system with r'_t being empty. Equivalently for the LP approaches, the impulse variable $\tilde{x}_t$ is the shock (or the innovation) itself. Consequently, no controls are needed in the OLS regression (1.6) to mop up any measurement error or serial correlation in the shocks, as typically done in many empirical applications (Ramey 2016, Stock and Watson 2018). As a result, the observed data vector $\bar{w}$ includes the shock itself as well as the macroeconomic variables of interest for both econometric models. The latter include: (i) output, (ii) consumption, (iii) investment and (iv) hours worked.

I estimate the dynamic response of each of these variables to one of the three most common aggregate shocks: (i) monetary, (ii) fiscal and (iii) technology shocks. For monetary and technology shocks using the innovation η_t or the shock ε_t will lead to identical estimated responses, however, this is not the case for the fiscal policy shock. Recall that in the Smets and Wouters model government spending is completely

exogenous but it is affected by the technology shock as follows:

$$(1.7) \quad \varepsilon_t^g = \rho_g \varepsilon_{t-1}^g + \rho_{ga} \eta_t^a + \eta_t^g$$

where ρ_{ga} captures the contemporaneous correlation between the two shocks. As a result, if $\rho_{ga} \neq 0$, using the shock ε_t^g without controlling for TFP will lead to incorrect responses. To circumvent this issue I will initially use the innovation rather than the shock itself as our impulse variable for all the shocks, i.e. $\tilde{x}_t = \eta_t^i$ for $i \in \{m, g, a\}$. Nonetheless, I will still explore the differences between using the innovation or the correlated fiscal policy shock as explained in Section 1.3.4.3.

1.3.4.2. *Recursive identification*

On the other extreme, I assume that the econometrician only observes the endogenous variables with no direct measure of the shock. Consistent with the large literature in recursive shock identification in VARs (e.g. see Christiano, Eichenbaum, and Evans, 1999), the shock of interest is the orthogonalized innovation to a policy variable i_t included in the vector of endogenous variables $\bar{w}_t$.

There are two common identification assumptions to impose recursive zero restrictions on contemporaneous coefficients (Ramey 2016). First, *the policy variable does not respond within the period to other endogenous variables*. For example, Blanchard and Perotti (2002) impose this constraint in the context of government spending shocks. And second, *other endogenous variables do not respond to the policy variable within the period*. Bernanke and Blinder (1992) were the first to identify monetary policy shocks in this way, but they have been followed by others like Christiano, Eichenbaum, and Evans (2005).

Consistent with this literature, we follow the second approach for monetary shocks and order the policy rate last as this restricts other variables in the VAR to not respond contemporaneously to the monetary innovations. Among the other macro variables in the VAR we include: (i) output, (ii) consumption, (iii) investment, (iv) hours worked, (v) wages, and (vi) inflation. On the other hand, we follow the first approach for the fiscal and technology shocks and use government expenditures or productivity as the first series in the VAR, respectively. For both shocks, we include: (i) output, (ii) consumption, (iii) investment, and (iv) hours worked as the other variables in the VAR.

Interestingly, in the context of the Smets and Wouters model these recursive assumptions will only be correct in the case of technology shocks as TFP is purely exogenous.

Government expenditures, despite being exogenous, are correlated with the productivity shock while real variables and prices respond contemporaneously to monetary innovations despite price and wage rigidities as shown in Figure 6 of Smets and Wouters' (2007) paper. These invalid identification assumptions for fiscal and monetary policy shocks make our estimation exercises more interesting as it will allow us to test to what extent *Ind. Inf.* is robust to this type of misspecification, i.e. one in which the recursively orthogonalized innovations do not equal the structural shocks in the Smets and Wouters model.

Turning to the LP responses, we know that any SVAR identification scheme can be also implemented using LP methods (Plagborg-Møller and Wolf 2021). In fact, for the identification strategy used for the technology and fiscal policy shocks, this only requires to set the impulse variable $\tilde{x}_t$ to the policy variable i_t ; while for the monetary policy identification scheme, we also need to control for the contemporaneous variables that are ordered before the policy variable in the VAR system.

1.3.4.3. *Noisy direct measures of the shocks of interest*

In between these two extremes, there is a growing and very popular strand of the literature that relies on external information to construct a direct measure of the shock of interest. These directly measured shocks often capture only part of shock or are measured with error (Stock and Watson 2018). For example, Romer and Romer (2004) use narrative methods to construct a monetary policy shock in which Greenbook forecasts are used to separate the Fed's superior information from the exogenous shock. Nonetheless, they still use additional recursive assumptions when studying the responses of output and prices as they do not view their shock as pure (Ramey 2016).² Consequently, I consider a third identification strategy in which the observed innovation / shock is contaminated with measurement error. In particular, I assume that the econometrician observes a proxy for the innovation of the shock:

$$(1.8) \quad \eta_t^{obs} = \eta_t + \sigma_v v_t$$

where v_t is an iid innovation with zero mean and a standard deviation of one. The IRF estimation approach is identical to the observed innovation case in Section 1.3.4.1 but replacing η_t by equation (1.8) and assuming that $\sigma_v = 0.5$.

² Other examples of this approach include the fiscal policy shock measure in Ramey (2011) which uses Business Week's articles.

In addition to the classical measurement error scenario, I also consider the possibility that the measured shock is correlated with other shocks, which would violate the exogeneity condition. Recall that this is the case for the government spending shock within the Smets and Wouters (2007) model – see equation (1.7). Hence, I compare the estimation results from targeting responses to fiscal policy which have been estimated using information about the innovation η_t^g versus those that rely on the actual correlated shock ε_t^g . Again, the estimation procedure is identical to the observed shock case, but with a different information set.

Finally, I consider the case in which the IRFs have been normalized using the unit effect of Stock and Watson (2018) as they show that fixing the shock units via normalization allows to capture the dynamic causal effect even in the presence of measurement error.

1.4. Design, Implementation & Evaluation

This section describes how I set up the Monte Carlo study to analyze the small sample properties of the *IRF matching* and the *Ind. Inf.* estimators that use VARs or LPs as the source of data moments or as the binding function, respectively. The analysis is based on the economic model described in Section 1.2 under the hypothesis that the DGP and the estimated model are the same.³ That is, the log linearized version of the Smets and Wouters (2007) model is used to generate time series of macroeconomic variables as well as time series for the innovation of the shocks. These series are then used to estimate impulse response functions using the econometric models described in Section 1.3.3 and under the different identification schemes explained in Section 1.3.4. Finally, these estimated responses, which summarize the dynamics of the model/data, are used as moments/targets in estimation to pin down the structural parameters.

The Smets and Wouters (2007) model has 36 structural parameters but to reduce the computational burden I focus on 8 of these: the intertemporal elasticity of substitution $\{\sigma_c\}$, the consumption habit parameter $\{h_c\}$, the elasticity of labor supply $\{\sigma_l\}$, the investment adjustment cost parameter $\{\varphi\}$, and the non-adjustment probabilities and indexation parameters for wages $\{\xi_w, \iota_w\}$ and prices $\{\xi_p, \iota_p\}$. The “true” values of these structural parameters are listed in Table 1.1, while the remaining ones are fixed at the mean estimated values by Smets and Wouters (2007) – see Table 1A & 1B in their paper.

³ I do not consider the alternative that the model is misspecified. For a statistical criteria that detects and identifies misspecification in structural models see Inoue, Kuo, and Rossi (2020). For the performance of simulated method of moments in the context of model misspecification see Ruge-Murcia (2007).

TABLE 1.1. True values of structural parameters

Parameter	Value	Interpretation
σ_c	1.38	Intertemporal elasticity of substitution
h_c	0.71	Habit parameter
σ_l	1.83	Elasticity of labor supply
φ	5.74	Investment adjustment cost parameter
ξ_w	0.70	Probability of non-adjustment (wages)
ξ_p	0.66	Probability of non-adjustment (prices)
ι_w	0.58	Wage indexation parameter
ι_p	0.24	Price indexation parameter

NOTE. This table depicts the true value of the estimated parameters from the Smets and Wouters model. Their values coincide with the mean estimates from their 2007 paper.

Using these parameter values, the “true” model is simulated $S = 100$ times for $T = 300$ periods.⁴ This artificial dataset is used to estimate the dynamic response of four macro aggregates: output, consumption, investment and hours worked to either monetary, fiscal or technology shocks over five years ($H = 20$ quarters). Hence, for each Monte Carlo draw and each estimation setup we target $84 = 21 \times 4$ moments.⁵ Note that the Monte Carlo distribution of these targets/data moments is identical for both *IRF matching* and *Ind. Inf.* exercises – it represents β in problems (1.2) and (1.3). Recall that for *Ind. Inf.* approach the estimated IRFs are also computed based on the model simulated data at each candidate parameter vector, $\beta(\Theta)$. In that case, the sample size is inflated by a factor $\tau = 10$. In theory, we know that the asymptotic distribution of the estimates depends on this choice as simulation uncertainty decreases when the length of the simulated series to the sample size increases. However, in practice, having very long simulated series increases the computational cost and is not needed to obtain accurate estimates. Ruge-Murcia (2012) shows how this choice affects the parameter estimates in the context of DSGE models estimated by SMM. Consequently, I do not explore this dimension and simply set this hyper-parameter to a common value used in practice. Nonetheless, I consider the case in which the data moments/targets are estimated on a sample with just $T = 100$ periods to study the *small sample* bias in LPs documented in Herbst and Johannsen (2023). For simplicity, this robustness test is

⁴ A sufficiently long burning sample is used to get rid of the initial conditions.

⁵ These IRF are chosen based on their sensitivity to changes in the estimated parameters, although one could have use a more formal selection criteria, e.g. the one developed in Hall, Inoue, Nason, and Rossi (2012). But, in either case, the selected targeted responses should not influence the results as long as they exhibit the bias-variance trade off documented in Li, Plagborg-Møller, and Wolf (2023).

performed only in the context of the observed innovation scheme, but also considers VARs estimated responses since they are not exempt from this type of bias.

Finally, I use the identity matrix as the weighting matrix $W = I$ since it is one of the most widely used in empirical work. Nevertheless, I also consider: (i) the inverse of the variance-covariance matrix of the data moments (VCM) as it is the optimal weighting matrix, and (ii) a diagonal matrix whose entries are the inverse of the IRFs horizon $1/h$. The latter tries to address the possible identification problem arising from the little and noisy information contained in impulse responses at long horizons (Canova and Sala 2009).

1.4.1. Performance Metrics

To evaluate the performance of a given estimator $\hat{\Theta}$ of Θ , we consider different metrics that can be classified into two groups: (i) *overall performance* metrics that speak about the structural estimation as a whole and consequently inform us about how the estimated model fits the DGP, and (ii) *parameter-by-parameter* metrics that look at each estimated parameter individually. Most of the literature focus only on the latter and assesses the performance of the estimation based on the bias and standard deviation of each estimated parameter and even sometimes on the sum of the two squared: the Root Mean Squared Error (Smith 1993, Ruge-Murcia 2007, 2012, 2020, Scalone 2018). Equally important is the overall fit, and hence, I stress the importance of these metrics in the discussion of our results as they sometimes draw a different picture.

1.4.1.1. Overall performance

The most natural metric that speaks about the overall performance of the estimation is the value of the objective function that one is trying to minimize, that is J^{smm} and J^{irf} in equations (1.2) and (1.3). These are often refer to as the J -statistic. A recurrent problem with this statistic is that it depends on the units of the weighting matrix W . Consequently, when reporting the value of the J -statistic for the different estimation setups we will fix the weighting matrix to the identity independently of which weights have been used during the optimization stage.

The J -statistic is frequently used in practice because it is easy to compute, however, it only gives an approximate sense of how well the estimated model is able to capture the dynamic responses to various shocks. Given that we control the DGP, one can do better by looking at the distance between the structural IRFs at the true parameter vector Θ^* and at the estimated one $\hat{\Theta}$ as shown below

$$(1.9) \quad J^* = \left(IRF(\Theta^*) - IRF(\hat{\Theta}) \right)' \left(IRF(\Theta^*) - IRF(\hat{\Theta}) \right)$$

Equation (1.9) can be computed for the targeted responses, but also for untargeted ones, e.g. output response to a technology shock when targeting monetary policy responses.

1.4.1.2. *Parameter-by-parameter performance*

The literature on DSGE estimation has looked at bias and standard deviations of the estimated parameters when evaluating different methods for obvious reasons. I also look at these metrics but with a small twist motivated by the loss function in Li, Plagborg-Møller, and Wolf (2023). Given the bias-variance trade off in estimated IRFs, and also, the heterogenous researcher's preferences about biases and noise in their parameter estimates, I consider a linear combination of bias and variance with different bias weights as shown below

$$(1.10) \quad \mathcal{L}_\omega(\hat{\Theta}_i, \Theta_i^*) = \underbrace{\omega \times \left(\mathbb{E} \left[\hat{\Theta}_i \right] - \Theta_i^* \right)^2}_{\text{bias}} + (1 - \omega) \times \underbrace{\text{Var}(\hat{\Theta}_i)}_{\text{variance}}$$

Note that for $\omega = 1$, the researcher is only concerned about bias. For $\omega \in (0.5, 1)$ the researcher is more concerned about (squared) bias than variance, while for equal weights $\omega = 0.5$, this metric is proportional to the mean squared error (MSE).

Then, when comparing two different approaches, for example one that uses VARs and other that uses LPs, I will compute the difference between the two loss functions for different bias weights and as a fraction of the true structural parameter value to make deviations comparable across parameters. That is, my preferred measure of parameter-by-parameter performance has the following form

$$(1.11) \quad z_i \equiv \frac{\left(\mathcal{L}_\omega(\hat{\Theta}_i^{VAR}, \Theta_i^*) - \mathcal{L}_\omega(\hat{\Theta}_i^{LP}, \Theta_i^*) \right)}{\Theta_i}$$

where $\hat{\Theta}_i^{VAR}$ and $\hat{\Theta}_i^{LP}$ denote the parameters estimated when using VAR or LP as the (auxiliary) econometric model, respectively.

1.5. Results

1.5.1. The Best Case Scenario: Observed Innovations as Benchmark

I start by discussing the Monte-Carlo results under the assumption that the econometrician observes the true innovation. This is a situation that would never occur in practice, however, it is a good benchmark to initially test the properties of the *Ind. Inf.* and *IRF matching* estimators. The targeted estimated responses under such assumption are depicted in Appendix A.2.1, where I show that there is a bias-variance trade off between LPs and SVARs in the context of the Smets and Wouters model. But, what are the implications of this trade off for the estimated structural parameters?

Table 1.2 shows the overall performance metrics for the two estimation strategies and econometric models while averaging across the three sources of variation and the four lag lengths considered. The sample size is set to $T = 300$ observations. Focusing only on the top block that relies on the identity as the weighting matrix for now, one sees that using LP responses as targets in an *IRF matching* exercise is a better idea (lower J^*) because their smaller bias. However, this is no longer true in an *Ind. Inf.* exercise where the SVAR approach is slightly better given that SVAR responses have lower variance and their larger bias is irrelevant for the *Ind. Inf.* approach as it is robust to this type of misspecification.

TABLE 1.2. Overall performance using the observed innovation

	IRF matching				Indirect Inference			
	J_{irf}	J^*	Time	J^*_{unt}	J_{smm}	J^*	Time	J^*_{unt}
Identity Matrix								
<i>Local Projection</i>	35.10	0.27	3.49 min	18.70	32.54	0.39	42.88 min	17.91
<i>Structural VAR</i>	35.23	0.41	3.93 min	17.93	33.87	0.33	14.47 min	18.39
Diagonal Matrix								
<i>Local Projection</i>	34.44	0.22	3.61 min	18.87	32.82	0.35	40.56 min	18.22
<i>Structural VAR</i>	34.87	0.27	3.85 min	18.20	34.17	0.31	11.55 min	18.62
Optimal Weighting Matrix								
<i>Local Projection</i>	33.63	0.04	3.07 min	21.56	32.69	0.06	35.56 min	21.41
<i>Structural VAR</i>	34.17	0.05	3.20 min	20.80	34.26	0.08	10.69 min	20.90

NOTE. This table shows the overall performance metrics and the average computing time for *IRF matching* and *Ind. Inf.* exercises that use either the identity, the diagonal or the optimal weighting matrix.

In terms of parameter by parameter performance, what seems to drive these differences between targeting the LP versus the SVAR estimated responses in an *IRF matching* exercise is the lower bias obtained for the inter-temporal and intra-temporal elasticities of substitution $\{\hat{\sigma}_c, \hat{\sigma}_l\}$, as shown in panel A of Figure 1.1 by the darker red color at $\omega \approx 1$. On the other hand, the better overall performance of the SVAR approach in the *Ind. Inf.* exercise is driven by the lower variance of the intra-temporal elasticity, the habit parameter, the investment adjustment cost and especially the Calvo (1983) probability of wage adjustment $\{\hat{\sigma}_c, \hat{h}_c, \hat{\varphi}, \hat{\xi}_w\}$, as shown by the blue bars in panel B of Figure 1.1. More generally, it is interesting to observe that for most estimated parameters the LP approach tends to do better when the researcher gives a lot weight to the bias, while the SVAR approach is more desirable under low bias weights.

Lag Length. Understanding these previous results requires to dig deeper into what drives the differences in the estimated IRFs. The lag length is a natural choice as the trade off between LPs and VARs depends on it. As shown in Figure A7 in Appendix A.3, LP responses are independent of the lag length and SVAR responses approximately agree with them up to horizon p . It is only beyond horizon $h > p$ where they disagree substantially. In fact, as discussed in Li, Plagborg-Møller, and Wolf (2023), it is the more




FIGURE 1.1. Parameter-by-parameter performance

NOTE. This figure show our preferred measure of parameter-by-parameter performance, equation (1.11), for both *IRF matching* and *Ind. Inf.* estimation approaches under the identity weighting matrix. Here, for each parameter consider in the estimation, a red color indicates that the LP outperforms the SVAR approach, while the blue color highlights the opposite situation: SVAR better than LP.

restrictive way in which SVAR extrapolate long run responses from the first p sample auto-covariances that yields the lower variance at a higher bias. Nonetheless, when increasing the lag length the confidence intervals of the SVAR responses increase and become more alike to those of the LP, which is consistent with the latest “no free lunch” result in Olea et al. (2024). So again, what are the implications of these results on LP and SVAR estimates for the structural parameters when LP/SVAR estimated responses are used as the source of moments in a partial information DSGE estimation?

Table A1 breaks down the overall performance of the *IRF matching* and *Ind. Inf.* approaches by the choice of the lag length. A couple interesting observations arise. First, for the *IRF matching* the J^* from the SVAR gets closer and closer to the LP counterpart as the lag length increases. This is mostly driven by the REDUCTION of J^* associated to the lower bias of the SVAR responses at long horizons. In fact, median and confidence intervals of the targeted IRFs are almost identical when $p = 12$, and consequently, estimated parameters and J^* are very similar too. And second, for the *Ind. Inf.* exercise, which recall is robust to misspecification, the gap in J^* is also decreasing but because J^* in the SVAR approach INCREASES as the confidence intervals of the SVAR responses get wider. Overall, it seems that the DSGE modeler will be better off by matching tightly estimated responses, independently of their bias, while using an *Ind. Inf.* approach. However, this comes at higher computational cost as it requires model simulation and IRF estimation at each iteration. Consequently, for some models *IRF matching* may be more suitable and therefore targeting a well estimated response with low bias becomes crucial.

Sample Size. The presence of small sample bias can become an issue for *IRF matching* for obvious reasons, but it can also affect *Ind. Inf.* as long as it also affects the variance of the responses. Consequently, the choice of the sample size used to generate the data moments / targets is another relevant dimension to understand the differences between the estimation approaches studied in this paper. Hence, I repeat the Monte Carlo estimations using a smaller sample of $T = 100$ observations since this is the typical sample length encounter in most macroeconomic applications (Herbst and Johannsen 2023). As shown in Figure A8, small sample bias in LP responses is also present in the baseline sample with $T = 300$ observations, however, it becomes larger when I reduce the sample size.⁶ Hence, I consider two avenues: (i) I investigate whether *Ind.*

⁶ Recall that Plagborg-Møller and Wolf’s (2021) result about $LP(p)$ exactly agreeing with the structural responses is a population result, i.e. for very large T . Panel C & D in Figure A8 illustrate this point.

Inf. improves upon *IRF matching* when the small sample bias is more severe in the LP approach, and (ii) I study whether correcting for bias in the data moments / targets using bias correction terms improves the overall performance of the estimation.

Table A2 addresses these two questions. First, by comparing the LP approach under the two sample sizes one sees that *Ind. Inf.* improves upon *IRF matching* when the small sample bias becomes very large at $T = 100$. Nonetheless, the performance of the estimation under both approaches is worse as the variance of the targets / data moments increases, which can be seen by comparing Figures A1 and A9. The bias in the SVAR is not related to the sample size, but smaller samples also increase the variance. As result, the overall performance when using SVAR responses with $T = 100$ is also worse than when $T = 300$ observations are employed. And second, when I repeat the estimation using the bias corrected versions of the LP and SVAR, discussed in Section ??, one can see that correcting for small sample bias is very effective when estimating the model via *IRF matching*. In fact, the J^* is around 1.5 times smaller when bias correction terms are employed to generate the targets. Finally, bias correction in the auxiliary models is not as relevant for *Ind. Inf.* estimation.

Weighting Matrices. All the previous discussions were based on the identity matrix which is a common choice in practice given its simplicity. However, I also explore the choice of two alternative weighting matrices. First, a diagonal matrix that has $1/h$ as its diagonal elements and hence gives a lower weight to the responses at longer horizons. And second, the optimal weighing matrix, which is known to be the inverse of the VCM of the moments.

To understand how this particular choice affects the overall performance of the estimation, start by looking at the second block of Table 1.2, which shows the J^* under the diagonal matrix. It is not surprising that the differences in J^* 's between using LPs or SVARs shrinks (relative to the identity matrix). Recall that at short horizons SVARs and LP responses approximately agree and consequently putting more weight on these coefficients imply more similar outcomes for the estimation. In fact, this is particularly strong for the SVAR targets in the *IRF matching* approach as it discounts the importance of matching the biased long-run responses. Another interesting observation is that J^* 's are generally lower for both estimation approaches and econometric models. Turning now to the last block of Table 1.2, which shows the overall performance metrics when the optimal weighting matrix is used, J^* 's are significantly smaller and remarkably close to 0, which is the best possible outcome. Moreover, the use of these more efficient

weighting matrices reduces the computational time, with the optimal weighting matrix being the best option among the three.

The parameter by parameter performance under each of these three alternative weighting matrices can be seen in Appendix A.3.3. The main takeaway is that the mean estimates improve substantially when a more efficient weighting matrix is used. Surprisingly, the improvement in terms of standard deviations is not as large as initially expected.

1.5.2. The Good Old-Fashioned Days: Recursive Identification

I now turn to discuss the estimation set ups that assume that the econometrician does not observe the shock, but she is able to recover it using restrictions based on economic theory. The most widely used approach is to impose zero restrictions on contemporaneous coefficients. As discussed in Section 1.3.4.2, I will identify technology and fiscal policy shocks by assuming that the policy variable, TFP or government spending, does not respond within the period to other exogenous variables. Importantly, this assumption will not hold for the fiscal policy shock in the Smets and Wouters model because government spending is contemporaneously correlated with the technology shock. Hence, I will postpone that discussion to Section 1.5.3 where I address the problem of identifying shocks subject to measurement error and its implications for the structural parameters. The results from targeting technology shocks can be found below in Section 1.5.2.1. Regarding the monetary policy shock, I instead assume that the policy variable, the interest rate, does not affect other endogenous variables within the period. This assumption does not hold in the Smets and Wouters model either and so I explore what are the consequences of targeting responses to misspecified VAR/LP models in Section 1.5.2.2 below.

1.5.2.1. Technology shock

The responses to a technology shock recursively identified within the Smets and Wouters (2007) model are identical to those obtained by assuming that the econometrician observes the innovation/shock. Obviously, the recursive assumption is correct and hence one can recover the true shock via a Cholesky decomposition. Hence, the estimation results using the minimum distance approach will be identical under the two assumptions. The first block of Table 1.3 shows the overall performance metrics where one sees that the main lessons from Section 1.5.1 still apply when focusing only on technology

shocks. Another interesting observation concerns how the model captures the dynamic response to other shocks, which is measured by J_{unt}^* . Its relatively large values across estimation approaches and econometric models indicate that targeting the response to technology shocks is not a great idea as the estimated model will miss the dynamic responses to fiscal and monetary policy at the optimal parameter vector. Further, Figure 1.2 shows how LP and SVAR compare when individually focusing on each estimated parameter. Such comparison is informative about the contribution of each estimated parameter to the overall outcome. Actually, one can confirm by looking at panel A that the superior performance from targeting LP-IRFs in the *IRF matching* estimation comes from a more accurate estimation of the investment adjustment cost parameter $\{\hat{\varphi}\}$. Similarly, the SVAR approach to *Ind. Inf.* is also better than the LP approach because it does a better job in pinning down φ . Notice that even though the SVAR approach to *Ind. Inf.* is also better at identifying other parameters, such as the intra-temporal elasticity of substitution $\{\hat{\sigma}_c\}$, these are not so relevant for shaping the responses to technology innovations. In fact, σ_c is better identified when targeting the SVAR-IRFs in a *IRF matching* exercise despite its overall performance is worse than when targeting LP-IRFs.




FIGURE 1.2. Breakdown of Figure 1.1 by targeted shock – Technology

NOTE. This figure show our preferred measure of parameter-by-parameter performance, equation (1.11), for both *IRF matching* and *Ind. Inf.* estimation approaches under the identity weighting matrix. Here, for each parameter consider in the estimation, a red color indicates that the LP outperforms the SVAR approach, while the blue color highlights the opposite situation: SVAR better than LP.

1.5.2.2. Monetary policy shock

In the Smets and Wouters model a negative monetary policy shock has a positive impact in real activity at time $t = 0$ as shown by the dashed lines in Figure A1 or A3. On the other hand, ordering the policy rate last in the VAR and recovering the responses through a Cholesky decomposition implicitly assumes that monetary policy does not have a contemporaneous impact on other endogenous variables. Consequently, all the estimated responses, either via LP or SVAR, start at 0 when the monetary policy shock has been identified in this way. Obviously, this assumption is at odds with the model. Thus, differently from the technology shock, I now investigate what are the implications for the structural parameters of targeting these misspecified responses.

The middle and bottom blocks of Table 1.3 show the overall performance metrics when targeting responses to the observed or the recursively identified shocks, respectively. Focusing first on the observed shock, one sees that in line with the previous results, targeting LP-IRFs is a better approach when relying on *IRF matching*, while using a SVAR is better than LP as a binding function for *Ind. Inf.*, even though only by a small margin in this case. Additionally, and differently from the technology shock, targeting the responses to monetary policy shocks are a good idea in the context of

TABLE 1.3. Decomposition by the targeted shock

	IRF matching				Indirect Inference			
	J_{irf}	J^*	Time	J_{unt}^*	J_{smm}	J^*	Time	J_{unt}^*
Technology shocks								
<i>Local Projection</i>	1.05	0.67	2.87 min	37.30	0.70	0.84	42.41 min	35.92
<i>Structural VAR</i>	2.53	1.07	3.11 min	35.74	0.97	0.66	14.34 min	37.31
Observed monetary policy shock								
<i>Local Projection</i>	50.65	0.07	3.46 min	9.36	48.46	0.31	41.39 min	9.40
<i>Structural VAR</i>	54.07	0.11	4.38 min	9.26	53.60	0.30	14.65 min	9.44
Recursive monetary policy shock								
<i>Local Projection</i>	48.11	0.29	3.34 min	9.60	56.91	0.18	78.57 min	9.34
<i>Structural VAR</i>	47.09	0.34	3.78 min	9.31	58.70	0.12	11.44 min	9.34

NOTE. This table shows the overall performance metrics for *IRF matching* and *Ind. Inf.* when estimated responses to technology shocks (top block), observed monetary policy shocks (middle block) or recursive monetary policy shocks (bottom block) are being targeted. In all set-ups we are averaging the results across different lag lengths.

the Smets and Wouters model as one would also be able to capture the dynamics of technology and fiscal policy fairly well, as shown by the lower J_{unt}^* (relative to the results obtained by targeting the technology shock).

Turning now to the estimation set-ups that targets the misspecified responses to the recursive monetary policy shock, one can see that when *IRF matching* is the estimation approach, overall performance gets worse as the larger bias of estimated responses relative to the true structural IRFs gets reflected into the estimated structural parameters, independently of the econometric model employed. On the other hand, *Ind. Inf.* is robust to misspecification and in fact improves upon the observed shock case: J^* is lower in the bottom block than in the middle block. This may seem surprising initially, but it is explained by the lower variance of the responses to recursive shocks. Imposing a zero contemporaneous response reduces the bands of the estimated IRFs that are used as data moments, and hence, structural parameters are more tightly estimated. Further, within the recursive shock, the LP approach outperforms the SVAR approach in a *IRF matching* exercise while the opposite is true in the *Ind. Inf.* approach. But what parameters are responsible for these overall estimation outcomes?

Figure 1.3 compares the difference in parameter-by-parameter losses for various bias weights as shown in equation (1.11) when a estimated responses to a recursive monetary policy shock are used as targets / data moments. Starting by panel A in which




FIGURE 1.3. Parameter-by-Parameter Performance - Recursive Monetary Policy Shock

NOTE. This figure shows our preferred measure of parameter-by-parameter performance, equation (1.11), for both *IRF matching* and *Ind. Inf.* estimation approaches that target responses to a recursive monetary policy shock under the identity weighting matrix.

the *IRF matching* is considered, one sees that again φ plays an important role and is responsible for explaining the better performance of the LP approach. Note that the intra-temporal elasticity of substitution σ_l is also better pinned down by the LP in the *Ind. Inf.* approach but still LP underperforms. In fact, as shown in panel B, almost all the other structural parameters are better estimated with the SVAR as the auxiliary model, independently of the bias weight.

1.5.3. Direct Proxies for the Shocks: Measurement Error & Unit Effect Normalization

Finally, I present the results of the Monte Carlo analysis that assumes that the econometrician does not observe the true shock but a proxy for it. I distinguish three cases: (i) the proxy is contaminated with a white noise error and the econometrician is not aware of it, (ii) the proxy is contaminated with a term that is correlated with other shocks in the system / model and the econometrician also does not correct for it in any way or form, (iii) the proxy is contaminated with white noise error but the responses are normalized such that the error cancels out by means of the Stock and Watson (2018) unit effect normalization. The targeted moments used in the structural parameter estimation under each of these variants are depicted in Appendix A.2.3.

1.5.3.1. Classical measurement error in the innovation

Figure A4 shows that the presence of measurement error in the innovation leads to attenuation bias. It arises from the variance term in the denominator of the least squares estimator and hence it is common to both econometric models, LPs and VARs. Since neither of these two IRF estimators are robust to the presence of measurement error, then using LP or SVAR estimated responses during the structural estimation stage won't affect the results in any different way than in the observed shock case. Nonetheless, the bias associated to the presence of measurement error will still worsen the structural estimation outcome for both *IRF matching* and *Ind. Inf.* estimators. Because targeted responses are now biased towards zero, then those parameters that dampen the IRFs are selected as optimal. Note that the econometrician is not aware of the measurement error and hence uses the true innovation in the model for updating the model counterpart of the IRFs for each parameter vector considered. As a result, the simulated / structural IRFs at each candidate vector do not suffer from attenuation bias. Then, a potential solution that may improve the estimation outcome will be to estimate the variance of the white noise error that contaminates the innovation. As a result, the attenuation

bias in the model moments can be introduced through this parameter rather than by driving the structural parameters away from their true values. This extension is left for future work.

The top block of Table 1.4 shows that the overall performance metrics of the structural estimation that targets estimated responses to a technology shock that features uncorrelated measurement error and confirms the above intuition. As shown by the J^* , the estimation outcome is significantly worse for both LPs and VARs as well as for the *IRF matching* and *Ind. Inf.* estimators relative to the observed shock case. Nevertheless, the lessons from the observed shock case still apply. That is, targeting LP-IRFs is a good idea when resorting to *IRF matching* given their low bias, while using VARs for *Ind. Inf.* is a better choice given their lower variance.

These findings also apply to other sources of variation such as monetary or fiscal policy shocks in which the innovation is also observed with classical measurement error. These results are shown in the middle and bottom block of Table 1.4. One can see there how the J^* is significantly larger relative to the observed shock case for both *IRF matching* and *Ind. Inf.* estimators.

TABLE 1.4. Shock proxies and classical measurement error

	IRF matching				Indirect Inference			
	J_{irf}	J^*	Time	J_{unt}^*	J_{smm}	J^*	Time	J_{unt}^*
A technology shock proxy ($\eta_t^{a,obs}$)								
<i>Local Projection</i>	1.79	1.25	3.05 min	34.30	1.35	1.40	40.23 min	33.31
<i>Structural VAR</i>	3.41	1.70	2.80 min	33.47	1.70	1.18	13.74 min	34.39
A monetary policy proxy ($\eta_t^{m,obs}$)								
<i>Local Projection</i>	46.81	0.33	3.46 min	9.73	45.99	0.61	40.57 min	9.70
<i>Structural VAR</i>	48.07	0.35	3.72 min	9.43	49.42	0.71	12.51 min	9.77
A fiscal policy proxy ($\eta_t^{g,obs}$)								
<i>Local Projection</i>	48.05	0.05	4.23 min	8.21	42.52	0.19	45.60 min	7.47
<i>Structural VAR</i>	44.04	0.19	4.07 min	7.80	41.41	0.14	13.73 min	7.62

NOTE. This table shows the overall performance metrics for *IRF matching* and *Ind. Inf.* when the shock used to estimate IRFs has been contaminated with classical measurement error.

1.5.3.2. *Correlated measurement error: government spending and its correlation with technology*

Now I consider the case in which the observed shock is correlated with other shocks. Recall that in the Smets and Wouters model this is the case of government spending. Differently from the previous scenario here I assume that this correlation is known during the optimization stage. That is, ρ_{ga} is neither set to zero nor estimated, but instead fixed to its true value when updating the model moments during the estimation.

The targeted responses to a government spending shock are shown in Figure A5.⁷ The estimation results from targeting these responses are shown in the top block of Table 1.5. The J^* is again much larger than in the observed shock case or in the proxy measure with classical measurement error, and for both estimation approaches. Hence, as expected, correlated errors are a bigger issue than uncorrelated ones for structural parameter estimation. Surprisingly, *Ind. Inf.* is not more robust to this type of biases than *IRF matching*. Thus, differently from (misspecified) recursive shocks, there is not an advantage in using *Ind. Inf.* over the traditional *IRF matching* approach when IRFs are estimated using proxies of the shocks.

1.5.3.3. *Unit normalization: a 1% increase in the policy rate*

Stock and Watson (2018) has shown a way of dealing with measurement error in the proxy variables by estimating relative rather than absolute impulse responses. The unit effect normalization is shown at work in Figure A6 where I plot the responses to a monetary policy shock estimated with a contaminated proxy but whose responses have been normalized such that the policy rate increases by 1% upon impact. The first thing to notice is that the population (dotted line) and the structural (dashed line) responses coincide at all horizons and for all variables. Nonetheless, there are still some discrepancies in finite samples as it was the case for the observed shock identification scheme. But how does this rescaling of the IRFs affect the structural parameters and the overall performance of the structural estimation?

The bottom block of Table 1.5 shows that the J^* is still larger than in the observed shock case, but the improvement upon the unnormalized responses is substantial. For example, the J^* coming from the *IRF matching* that targets LP-IRFs is 0.07 and 0.08 in the observed shock and normalized responses, respectively; while it equals 0.33 when using

⁷ These responses are identical to those obtained when ordering government spending first in the VAR and inferring the shock recursively.

TABLE 1.5. Correlated shocks & unit normalization

	IRF matching				Indirect Inference			
	J_{irf}	J^*	Time	J_{unt}^*	J_{smm}	J^*	Time	J_{unt}^*
A correlated fiscal policy proxy ($\varepsilon_t^{g,obs}$)								
<i>Local Projection</i>	30.82	0.34	4.09 min	7.80	39.05	0.35	46.13 min	10.15
<i>Structural VAR</i>	31.45	0.34	4.19 min	7.78	42.42	0.40	14.20 min	10.53
A 1% increase in r_0 (Stock and Watson (2018) normalization)								
<i>Local Projection</i>	50.77	0.08	3.83 min	19.34	49.49	0.52	49.84 min	17.85
<i>Structural VAR</i>	53.41	0.32	4.04 min	18.86	51.23	0.42	12.49 min	17.93

NOTE. This table shows the overall performance metrics for *IRF matching* and *Ind. Inf.* when the shock used to estimate IRFs has been contaminated with measurement error.

the responses to the monetary policy shock contaminated with classical measurement error. As the unit normalization helps in correcting the bias in estimated responses, it is very effective when employed in an *IRF matching* exercise. For *Ind. Inf.* the bias is less relevant and consequently the unit effect normalization is not as effective. In fact, the J^* for the VAR is 0.42 when using the normalization, 0.71 without normalization and 0.30 in the observed shock case. Finally, here the main lesson from the observed shock still applies and using LP-IRFs is better for *IRF matching* while SVAR-IRFs are more effective in *Ind. Inf.* estimations.

1.6. Conclusion

This paper conducts a Monte Carlo analysis to examine the small sample performance of *IRF matching* and *Ind. Inf.* estimators that target IRFs that have been estimated with LP or VAR models. I drew the following five conclusions:

1. The bias-variance trade off between LP and SVAR estimated IRFs affects the estimated structural parameters obtained with minimum distance estimators such as *Ind. Inf.* and *IRF matching* estimators. Nonetheless, it affects them differently. *IRF matching* is more sensitive to bias in targeted responses and hence using LP-IRFs is preferable, while *Ind. Inf.* is robust to misspecification and hence benefits from the lower variance of VAR-IRFs.
2. The number of lags used in the VAR or as controls in the LP is crucial in understanding not only the differences between estimated IRFs but also in the

estimated structural parameters. When the lag length p is large, then IRFs and estimated parameters are similar independently of the econometric model used. On the other hand, when p is small LP-IRFs are less biased and hence better for *IRF matching*, while SVAR-IRFs have a larger bias but lower variance which helps when estimating the parameters via *Ind. Inf.* as the later is robust to these type of biases in estimated responses.

3. The small sample bias of LPs, documented by Herbst and Johannsen (2023), worsens the performance of the structural estimation, specially in the case of *IRF matching*. Using their bias correction term for the targeted moments improves the estimation outcome of the *IRF matching* estimators, while it is irrelevant for *Ind. Inf.* applications.
4. Incorrect recursive identification for the target moments are not an issue for the estimation of structural parameters as long as *Ind. Inf.* is employed. However, it is problematic for *IRF matching*.
5. The presence of measurement error in the proxies used to estimate IRFs worsens the structural estimation outcome for both estimation methods and econometric models. Using the unit effect normalization of Stock and Watson (2018) help ameliorating this problem.

These findings are applicable to a wide range of economic models as they rely on the statistical properties of the impulse response estimators (e.g. VARs suffering from truncation bias, while LPs being independent of the lag length) and those of the structural parameter estimators (e.g. *Ind. Inf.* being robust to misspecification). Moreover, the Smets and Wouters (2007) model contains many ingredients that are still present in many modern macro models. For example, the production side of the nowadays very popular HANK models shares many elements with the Smets and Wouters (2007) model. In this context, it is important to recognize still that one would need to understand how the additional heterogeneity would be reflected in the LP and VAR responses, and in turn, in the structural parameter estimates. Nevertheless, as long as this additional heterogeneity does not generate any non-linearities in the responses to shocks, the lessons learned in this paper should still apply. Recent developments in impulse response estimation also seem to point into the world of heterogeneity through the estimation of state-dependent responses. Minimum distance estimators like the ones considered in this paper are well suited to deal with this type of non-linearities, however, its presence

may alter some of the conclusions in this paper and hence is an interesting avenue for future research. In this direction, a good starting point is the work of Ruge-Murcia (2020), which studies limited information estimation using non-linear solution methods for the economic model and state-dependent LPs for the auxiliary model. In any case, further research is still needed to understand the implications of heterogeneity for the structural parameter estimates and an economy featuring non-linearities arising from the heterogeneous behavior of agents within the model would be the perfect laboratory for this purpose.

Another interesting extension would be to consider the Bayesian counterpart of IRF matching and simulated method of moments given that after the work of Christiano, Trabandt, and Walentin (2010) part of the literature has moved towards this approach (e.g. see Christiano, Eichenbaum, and Trabandt (2016) or Bianchi, Ilut, and Saijo (2024)).

Chapter 2

The Aggregate and Distributional Implications of Credit Shocks on Housing and Rental Markets

Joint with Andrew Hannon & Gonzalo Paz-Pardo

Abstract: We build a model of the aggregate housing and rental markets in which house prices and rents are determined endogenously. Households can choose their housing tenure status (renters, homeowners, or landlords) and the size of their homes depending on their age, income and wealth. We use our model to study the impact of changes in credit conditions on house prices, rents and household welfare. We analyse the introduction in Ireland in 2015 of policies that limited loan-to-value (LTV) and loan-to-income (LTI) ratios of newly originated mortgages and find that, consistently with empirical evidence, they mitigate house price growth but increase rents. Homeownership rates drop, and young and middle-income households are negatively affected by the reform. An unexpected permanent rise in real interest rates has similar effects – by making mortgages more expensive and alternative investments more attractive for landlords, it increases rents relative to house prices.

2.1. Introduction

Housing is the largest asset in most household portfolios and housing-related expenses, such as rental payments, represent a substantial share of households' consumption baskets (Piazzesi and Schneider 2016). Over the last decade, large increases in rental prices and concerns about housing affordability, particularly for the young, have brought housing and rental markets to the forefront of the political debate in many advanced economies. At the same time, many countries have introduced macro-prudential measures to prevent the buildup of excessive household leverage, potentially constraining the access of many first-time buyers to mortgage credit.

Because housing and rental markets are closely connected, understanding how credit shocks, housing policies or developments in housing supply affect households requires studying how they impact both markets. For example, supply constraints that restrict the building of new housing might generate increases in both housing prices and rents. On the other hand, a credit tightening might push prospective buyers into renting, thus decreasing house prices and increasing rents (Gete and Reher 2018).

In this paper we study how households are affected by a shock that reduces the availability of mortgage credit, both through its direct impact and its equilibrium effects on rents and house prices. To do so, we develop a model of the rental and housing markets with two key features. First, households are heterogeneous as they differ in their age, income, and wealth, and make endogenous housing tenure choices which lead them to be renters, homeowners, or landlords of homes of different sizes. To get on and climb the property ladder, households can borrow through long-term mortgages for which downpayment and other constraints only hold at origination. And second, both rental and housing markets must be in equilibrium, which implies that house prices and rental prices must adjust to clear both markets as a result of potential shocks, but they may do so in different directions.

This flexibility contrasts with standard assumptions in macroeconomic models with housing, in which the rental sector is either non-existent or is owned by a deep-pocketed risk-neutral investor, implying that rents are fixed to a constant fraction of house prices. By allowing households to choose to become landlords in the context of a standard savings and portfolio choice model, we endogenously generate a distribution of landlords who are heterogeneous in their income, wealth and real estate holdings, with many of

them being small owners.¹ As a result, the model displays an upward sloping rental supply curve at a given house price without assuming that the rental and owner-occupied markets are segmented (Greenwald and Guren 2024). Our framework with endogenous landlords is close to that in studies of the tax treatment of housing (Sommer and Sullivan 2018), but accounts for the lumpiness of housing and thus generates empirically reasonable rental supply elasticities (Rotberg and Steinberg 2024).

We use the model to show that restricting credit access to potential mortgagors increases rental prices, reduces house prices and decreases the homeownership rate. The intuition is that, in the presence of binding constraints to mortgage credit, prospective homeowners need to either: (i) postpone or cancel their home buying decisions and stay renters for longer or (ii) downsize and purchase a smaller house with a smaller mortgage. Option (i) implies that more landlords need to enter the market, buy housing and provide it for rent. Because the marginal prospective landlord must be compensated above the previous one in order to step in and provide additional rental housing, rents go up to clear the market at a given house price. Option (ii), downsizing, pushes house prices down as the share of low quality/small houses increases. However, given that housing is lumpy and that no houses are available below a certain minimum size, some households do not have access to the downsizing adjustment channel and must switch to renting. The relative strength of these two channels determines the relative changes of house prices and rents. The more households choose to buy smaller houses rather than become renters, the smaller the effect on rental prices and the larger the house price drop.

In our main experiment, we analyze a borrower-based macro-prudential policy intervention that imposes maximum loan-to-value (LTV) and loan-to-income (LTI) ratios to newly originated mortgages, focusing on the introduction in Ireland in 2015 of a minimum 20% downpayment on a house and a maximum 3.5 ratio of mortgage debt to household income. This intervention, which was largely unanticipated and binding for many prospective buyers, is an excellent case study for the credit shocks we model. Acharya, Bergant, Crosignani, Eisert, and McCann (2022) showed empirically that this policy led to a reduction in house price growth in the areas in which the limits were

¹ This feature corresponds to the structure of rental markets in many advanced economies. In Ireland, which is the object of our study, more than 50% of all tenancies are held by households who only own one or two rental properties (see Appendix B.1.1). French administrative data shows that most rental properties are held in lumpy quantities by undiversified, small and home-biased landlords (Levy 2022). Even in the United States, where institutional investors are quite developed, private individual investors own 71.6% of all rental properties, with 14 million of them owning between 1 and 4 properties (Pew Research 2021).

particularly binding. We extend their analysis to rental prices and find that, consistent with our model mechanism, the reform led to a larger acceleration in rental price growth in those areas in which it had stronger effects.

We use our model, calibrated to the Irish economy, to quantify the short and long-run effects of the reform while keeping all other features of the housing and rental markets fixed. We find that the rent-to-price ratio increases by 2.79% upon impact, reaches its maximum (a 3.66% increase) during the fifth year after the introduction of the reform and attains a level similar to that of the new equilibrium in 10 years. Most of these dynamics are explained by the evolution of the rental price, as the average house price drops, but very little - highlighting the importance of realistic rental supply elasticities. Moreover, the reform reduces the homeownership rate by 1.8 percentage points in the long run and generates an increase in wealth concentration as landlords hold a larger fraction of the aggregate housing stock.

Our model reveals that constraining housing credit is particularly harmful for renters, young households and those in the bottom and middle of the income distribution not only because they find it harder to obtain a mortgage and need to postpone their buying decision, but also because they face higher rental prices in equilibrium. Prospective landlords, older households and those in the top of the income distribution benefit because they obtain higher returns from their housing investments. Overall, the 25 year old agents that are impacted by the reform suffer a loss equivalent to 1% of their lifetime consumption. Our results provide a first measure of the unintended, but large heterogeneous costs imposed on households by macro-prudential policies through both the rental and housing sectors and point to a redistribution of resources from poor renters to richer landlords. Nonetheless, we cannot measure the benefits of those policies given that we do not model the cyclical buildup of risk in the financial sector and, as a result, we are silent on optimal LTV or LTI ratios.

In our second exercise, we study a permanent, exogenous rise in the real interest rate of 1 percentage point. Compared to the macro-prudential intervention, this shock not only affects new buyers, but also existing mortgagors with a floating-rate mortgage who see their payments go up, and savers in financial assets who see their returns increase. This shock also reduces the homeownership rate, increases rents and reduces the average house price. Differently to the macro-prudential experiment, many households react to the more expensive mortgage rate by buying comparatively smaller houses and thus acquiring smaller mortgages. Additionally, because the rise in the real interest rate on savings makes financial assets comparatively more attractive than housing

for prospective landlords, the rental rate raises further to keep the rental market in equilibrium. The welfare impact is also highly heterogeneous, with households at the bottom of the income distribution losing and those at the top benefitting. Although we model a permanent increase in the real interest rate, the implications we find are consistent with the empirical evidence for monetary policy shocks, such as for example that in Dias and Duarte (2019). Thus, our real interest rate experiment suggests that tightening cycles of monetary policy that raise real interest rates may benefit some households via the reduction in asset prices (including housing), but may also make it harder for low-income households to afford increasing rents.

Related Literature. A broad literature studies the role of credit in driving the boom and bust cycle in house prices that was associated with the Great Recession (Favilukis, Ludvigson, and Van Nieuwerburgh (2017), Greenwald (2018) or Justiniano, Primiceri, and Tambalotti (2019)), while other papers studying this period focus on the role of liquidity in housing markets (Garriga and Hedlund 2020) or house price expectations (Kaplan, Mitman, and Violante 2020). Recently, Greenwald and Guren (2024) show that the implications of these models depend on their assumptions about rental markets. In particular, in a model in which rental and owner-occupied properties are identical, an increase in the homeownership rate does not impact house prices as households buy these additional houses from deep-pocketed landlords that do not use credit. Instead, when markets for rental and owner-occupied housing are segmented, an increase in housing demand raises house prices.

Our paper builds precisely on the intuition that modeling the rental market is key to understand house price dynamics. Differently from Greenwald and Guren's (2024) economy, we do not assume that the market between rental and owner-occupied properties is segmented. In our model, instead, the rent-to-price ratio of housing moves in response to a credit shock because of two reasons. First, even with a single house type, the supply curve for rental accommodation at a given house price is upward sloping due to the heterogeneity in landlords' stochastic discount factors. Although house prices are still the expected discounted future value of rents as in the standard user cost formula (Poterba 1984), the marginal landlord who is pricing rental housing can change endogenously, and thus credit shocks that reduce the homeownership rate can push rents upwards. Second, with heterogeneity in housing qualities, households react to the shock by moving to smaller properties, either owned or rented, which pushes house prices down in equilibrium. Both forces act together moving rent-to-price ratios

up. The ability to separately study movements in house prices and rents, rather than just looking at their ratio, is an additional contribution of our framework.

Our endogenous landlords are closest to those in studies of the tax treatment of housing (Floetotto, Kirker, and Stroebel (2016) or Sommer and Sullivan (2018)). Because of our focus on credit shocks, we introduce long-term mortgages, which allow us to study the effects of LTV and LTI ratios which only bind at mortgage origination. Additionally, in these frameworks homeowners decide every period which share of their home they rent out, while in ours households become landlords by buying additional discrete housing units. This more realistic feature introduces an additional friction to rental supply through the lumpiness of housing, and implies reduced rental supply elasticities which are much closer to those that we observe in the data. As Rotberg and Steinberg (2024) show in recent work, replicating these elasticities is key to understand the responses of rents to shocks or policies such as the mortgage interest deduction. Compared to their model, we do not explicitly target the elasticity by modeling the supply curve of a rental company, but let prospective landlords determine it endogenously.

This endogenous determination of rents and house prices, although novel to macro applications with heterogeneous households that study credit shocks, is also present in state-of-the-art equilibrium models of local housing markets, such as Favilukis, Mabile, and Van Nieuwerburgh (2022). On the household side, our model builds on the partial equilibrium framework in Paz-Pardo (2024), but it is augmented to allow households to own multiple properties and lease them out in the rental market.

We use our model to study the effects of borrower-based macroprudential interventions. Recent empirical literature has shown that the introduction of LTV and LTI limits reduces mortgage leverage (Van Bakkum, Irani, Gabarro, and Peydró 2023) and cools down tensioned housing markets (Acharya et al. 2022). We contribute to this work by showing that it increased rental prices, which is consistent with Gete and Reher (2018), who find that rents increased as a result of the contraction of mortgage supply in the United States after the Great Recession.

Our results in terms of the costs imposed by mortgage regulation on heterogeneous households complement a broad theoretical literature that shows that macroprudential frameworks are useful to guarantee financial and macroeconomic stability, like Laminetti, Mendicino, and Punzi (2013) or Farhi and Werning (2016). In recent work, Ferrero, Harrison, and Nelson (2023) and Muñoz and Smets (2022) focus on countercyclical borrower-based macroprudential rules and show how these interact with either monetary policy or credit to the large institutional investors in the rental market, respectively.

Oliveira and Queiró (2023) study the effects of the LTV and Payment-to-Income (PTI) constraints implemented in Portugal in 2018 in a framework based on Kaplan, Mitman, and Violante (2020) and find that the reform is welfare reducing due to changes in homeownership and the quality of housing.

The study of the link between monetary policy and house prices has a longer history. Iacoviello (2005) introduces housing in a business cycle model and finds that, although house prices react to monetary policy, there are little gains for the monetary authority to react to asset prices. Aastveit and Anundsen (2022) show that an expansionary 1 percentage point monetary policy shock raises house prices between 3 and 7 percent, depending on local housing supply elasticities. Dias and Duarte (2022) highlight, like we do, that monetary policy shocks increase the demand for renting with respect to home-owning, and as a result, rents tend to rise. However, in their model the changes in rent-to-price ratios are driven by the different relative stickiness of prices and rents, rather than through the endogenous formation of new landlords. Amaral, Dohmen, Kohl, and Schularick (2024) study the effect of a persistent decline in the real interest rate across geographical areas and highlight that it can have different impacts on rents and house prices depending on the initial rent to price ratio in a given location.

Overview. The rest of the paper is structured as follows. In Section 2.2, we present the model. In Section 2.3, we analyze the Irish macroprudential reform of 2015. First, by presenting some empirical evidence in section 2.3.1; and then by using a calibrated version of the model to analyze the effects on quantities and prices as well as on welfare. We discuss the parametrization in Section 2.3.2, and present model results in Section 2.3.3. In Section 2.4 we study the effects of a permanent real rate increase. Finally, Section 2.5 concludes.

2.2. The Model Economy

Our model economy is populated by households that differ in their age, income and wealth. They supply labor inelastically to a competitive production sector during their working age and make decisions about non-durable consumption, savings and housing tenure. Thus, they choose endogenously whether they are renters, homeowners or landlords. Although owning a house provides higher utility than renting, some households are forced or choose to rent because of binding credit constraints and up-keeping costs. At the other end of the spectrum, there are some households that own more than one house to lease them out and earn extra income. These heterogeneous renters and

landlords meet in a competitive market and determine the equilibrium rental rate. The housing stock is built by a construction sector that uses land permits and structures. The latter are produced, together with the final consumption good, by a competitive firm that uses labor as its only factor of production. The final good's price acts as the *numeraire*, while the house price is determined by the intersection between the supply from the construction sector and the demand from households.

2.2.1. Production

Final-good sector. The final-good sector operates a linear technology

$$(2.1) \quad Y_c = A_c N$$

where A_c is the constant aggregate labor productivity and N are the units of labor services. These firms hire labour in a competitive labour market, which implies that their profit maximization yields an equilibrium wage $w = A_c$. Final goods, whose price is normalized to 1, can be used both for household consumption, C , or as an intermediate input for the production of the housing good, in which case we label them structures, S . That is,

$$(2.2) \quad Y_c = C + S.$$

Construction sector. The construction sector operates a Cobb-Douglas technology

$$(2.3) \quad Y_h = A_h L^{\alpha_L} S^{1-\alpha_L}$$

where S is the quantity of structures, L is the amount of buildable land or housing permits in a given period, $\alpha \in (0, 1)$ is the constant share of land in production, and Y_h is the quantity of the housing good produced. We assume that the total amount of housing permits every period is fixed and they are priced competitively. Hence, the housing developer solves the following static problem

$$(2.4) \quad \max_{S,L} p_h A_h L^{\alpha_L} S^{1-\alpha_L} - p_L L - S$$

where p_L is the equilibrium price of buildable land. The first order conditions of the competitive housing developers' problem imply the following relation between housing

production Y_h and the house price:

$$(2.5) \quad Y_h = A_h^{1/\alpha_L} ((1 - \alpha_L) p_h)^{(1-\alpha_L)/\alpha_L} \bar{L}$$

where $\bar{L}$ is the aggregate amount of housing permits every period. Consequently, the elasticity of aggregate housing supply to house prices is:

$$(2.6) \quad \epsilon_{Y_h, p_h} \equiv \frac{\partial Y_h}{\partial p_h} \frac{p_h}{Y_h} = \frac{(1 - \alpha_L)}{\alpha_L} .$$

Housing comes in different qualities that represent different bundles or aggregations of the housing good. We denote them as $\tilde{h} = \{\tilde{h}_1, \dots, \tilde{h}_N\}$, and make the assumption that the final transaction price for each of these types is a multiple of the per-unit housing price. That is, $p(\tilde{h}) = \tilde{h} p_h$.

The construction sector's output is used for three purposes: the production of new houses, the upkeep of existing houses and the costly conversion between housing types. Upkeep costs are both for regular maintenance, amounting to δ_h per unit of housing in every period, and for the refurbishment of houses occupied by a terminal-age household after the occupant dies. The latter force implies that in every period $1/J$ of the housing stock needs to be rebuilt, where $1/J$ is the population share of terminal-age households. With respect to conversion costs, we assume that each unit of housing which is converted from one housing quality into another incurs a proportional cost ξ .

The aggregate housing stock, i.e. the aggregate amount of the housing good, can be measured as the quality-weighted sum of all housing units in the economy. If we let H_n^{sh} denote the share of houses of quality n ,

$$(2.7) \quad H = \sum_{n=1}^N H_n^{sh} \tilde{h}_n .$$

As a result, the law of motion for the housing stock is akin to a standard capital accumulation equation with the presence of adjustment costs:

$$(2.8) \quad H_{t+1} = \left(1 - \delta_h - \frac{1}{J}\right) H_t + Y_{h,t} + \sum_{n=1}^N \mathbb{1} \left\{ H_{n,t+1}^{sh} - H_{n,t}^{sh} < 0 \right\} \xi \left(H_{n,t+1}^{sh} - H_{n,t}^{sh} \right) H_t .$$

2.2.2. Households

Demographics. Household's age is indexed by $j = 1, \dots, J$. In the first $J^{ret} - 1$ periods they work. Thereafter they are retired until they die with certainty at age $J + 1$.

Preferences. Households derive utility from non-durable consumption and housing services. They value these streams of consumption according to

$$(2.9) \quad \mathbb{E}_0 \left\{ \sum_{j=1}^J \beta^{j-1} \frac{\left(c_j f(h_j, \tilde{h}_j) \right)^{1-\gamma}}{1-\gamma} \right\}$$

where $\beta \in (0, 1)$ is the discount factor, $\gamma > 0$ captures both risk aversion and intertemporal elasticity of substitution, $c > 0$ is consumption of non-durables, and f is a function of the number of houses owned h_j and the housing quality of the house in which the household lives $\tilde{h}_j$ and it is given by

$$(2.10) \quad f(h, \tilde{h}_j) = \begin{cases} \left(\tilde{h}_j / \tilde{h}_1 \right)^{\alpha_h} & \text{if } h = 0 \\ \left(\left(\tilde{h}_j / \tilde{h}_1 \right) \theta \right)^{\alpha_h} & \text{if } h \geq 1 \end{cases}$$

and as standard in the literature reflects that the housing service flow for homeowners is larger than for renters as reflected by $\theta > 1$, as well as the larger utility flow from better quality housing, captured by the ratio of the housing quality j to the lowest quality available $\tilde{h}_1$.

Endowments. Working-age households receive an idiosyncratic labor income endowment. We assume that it has a deterministic component that depends on age and a stochastic, persistent component. That is

$$(2.11) \quad \log y = \log A_c + g(j) + \eta$$

where A_c is an index of aggregate productivity, $g(j)$ is a polynomial in age and η represents the stochastic persistent component of earnings. We estimate the earnings process non-linearly as in De Nardi, Fella, and Paz-Pardo (2020) – see Section 2.3.2.1 for details. Retired households receive a fixed fraction of their last working period income for the rest of their lifetime. Households are also born with an initial endowment of liquid wealth that is drawn from a log-normal distribution. We also assume that they start their life as renters and thus have no housing wealth.

Liquid assets. Households can save in a one-period risk-free bond, $a \geq 0$ that yields a constant interest rate $r_s = r$, which is determined in the world market and is therefore exogenous.

Housing choices. Households decide on the quantity h and the quality $\tilde{h}$ of the housing they acquire. Households that do not own a house ($h = 0$) must rent one in the market at a unit rental rate p_r . For simplicity, we assume that the quality of this rental unit is as good as the lowest available quality in the owner-occupied market. Owner-occupiers ($h = 1$) choose the quality of the house that they live in across all possible $\tilde{h}$. Therefore, when a homeowner buys additional houses as an investment ($h > 1$), she purchases houses of quality $\tilde{h}_1$, rents them out and receives p_r per period and per house.

For both homeowners and landlords, there are some costs associated with housing purchases beyond its transaction price. Housing is illiquid. Consequently, we assume that households pay a proportional transaction cost that depends on the value of the house being sold or bought, $\tau_h p(\tilde{h})$. This cost captures real estate agent fees, taxation and other administrative costs. Houses are also costly to maintain. Therefore, homeowners and landlords pay maintenance costs to keep up with their depreciation, $\delta_h p(\tilde{h})$, where δ_h is the housing depreciation rate. When there is a transaction in the housing market, these costs are covered by the seller.

Mortgages. The purchase of a house can be financed through a mortgage at a fixed rate $r_b = r(1 + \kappa)$, where $(1 + \kappa)$ is the intermediation wedge between the mortgage rate and the risk-free rate. To reduce the dimensionality of the household problem, we treat mortgages as negative asset holdings $a \leq 0$, which prevents mortgagors from simultaneously having liquid assets. In other words, a denotes the net asset position.

The borrower must satisfy two constraints. First, a maximum loan-to-value (LTV) limit, which imposes that the size of the mortgage has to be smaller than a fraction of the value of the house. And second, a loan-to-income (LTI) requirement that limits household's borrowing to a multiple of its current (annual) income. Formally,

$$(2.12) \quad a' \geq -\lambda_{LTV} p(\tilde{h}') h'$$

$$(2.13) \quad a' \geq -\lambda_{LTI} y$$

where λ_{LTV} and λ_{LTI} are parameters. These two constraints only hold at origination.²

² These constraints are potentially different for first-time buyers and for buy-to-let investors. We detail how in Appendix B.2.2.

After the mortgage contract is signed and the house is purchased, the borrower chooses the repayment schedule freely. We make this modeling choice, instead of allowing for mortgage default, because in Europe delinquency – contract being breached by underpaying – is more frequent than foreclosure – contract being terminated (Hannon 2023). Nonetheless, we impose that: (i) all debts must be paid before the terminal age J , i.e. $a_J = 0$ and (ii) interest payments and a minimum amortization payment must be made in each period. As in Kaplan, Mitman, and Violante (2020), the minimum payment is determined by the constant-amortization formula

$$(2.14) \quad m_j = \frac{r(1+\kappa)(1+r(1+\kappa))^{J-j}}{(1+r(1+\kappa))^{J-j} - 1}.$$

Optimization Problem. A household of age j , income y , with h houses of quality $\tilde{h}$ and a assets solves the following dynamic programming problem

$$(2.15) \quad \begin{aligned} V(a, \underbrace{\{h, \tilde{h}\}}_{=s}, y, j) &= \max_{c, a', s'} \left\{ \frac{(c f(s))^{1-\gamma}}{1-\gamma} + \sigma_\varepsilon \varepsilon(s) + \beta \mathbb{E} V(a', s', y', j+1) \right\} \\ &\text{s.t.} \\ c + a' + p(\tilde{h}')h' + \mathbb{1}_{sell} \tau_h p(\tilde{h})h + \mathbb{1}_{buy} \tau_h p(\tilde{h}')h' + \delta_h p(\tilde{h})h &\leq \\ y + (1+r(1+\mathbb{1}_{a'<0}\kappa))a + p(\tilde{h})h + p_r(h-1) & \\ a' \geq \begin{cases} \max \left\{ -\lambda_{LTV} p(\tilde{h}')h', -\lambda_{LTI} y \right\} & \text{if } h' > h \\ a(1+r(1+\kappa) - m(j)) & \text{if } h > 0 \text{ and } a < 0 \\ 0 & \text{otherwise} \end{cases} & \end{aligned}$$

where $\sigma_\varepsilon \varepsilon(s)$ are choice-specific random taste shocks that are *i.i.d.* Extreme Value Type I distributed with scale parameter σ_ε . These represent shocks to the utility of homeownership (i.e., they are alike to moving shocks), but are also computationally convenient as they help to smooth out expected value functions (Iskhakov, Jørgensen, Rust, and Schjerning 2017).

2.2.3. Equilibrium

For a given risk free rate r , a competitive equilibrium in this economy consists of: (i) a value function, a housing choice probability, and a consumption and asset policy func-

tion for the households: $\{V, \mathbb{P}(h, \tilde{h}), c, a'\}$, (ii) a stationary distribution over households' state: $\{\mathcal{D}\}$, (iii) policy functions for the firms: $\{N, L, S\}$, and (iv) prices: $\{w, p_L, p_h, p_r\}$ such that they jointly solve the household, final-good firm and construction firm problems, as well as satisfy the following market clearing conditions:

$$(2.16) \quad \sum_{h=0}^X (h-1) \left(\int \int \sum_{j=1}^J \mathcal{D}(a, s, y, j) da dy \right) = 0$$

$$(2.17) \quad Y_h = \left(\delta_h + \frac{1}{J} \right) \underbrace{\sum_{n=1}^N \tilde{h}_n H_n^{sh}}_{=H} - \sum_{n=1}^N \mathbb{1} \left\{ H_{n,t+1}^{sh} - H_{n,t}^{sh} < 0 \right\} \xi \left(H_{n,t+1}^{sh} - H_{n,t}^{sh} \right) H_t$$

$$(2.18) \quad Y_c = C + S$$

$$(2.19) \quad \bar{L} = L$$

where equation (2.16) guarantees the equilibrium in the rental market, i.e. the demand for rental units by renters ($h = 0$) equals the supply of rental units by landlords ($h > 1$). Meanwhile, equations (2.17), (2.18) and (2.19) ensure that the housing, the goods and the land permits market clear. Note that the equilibrium per-unit house price p_h can be recovered analytically by substituting the housing investment function (2.5) into market clearing condition (2.17), which results in

$$(2.20) \quad p_h = \frac{1}{1 - \alpha_L} \left(\frac{1}{A_h} \right)^{\frac{1}{1 - \alpha_L}} \left(\frac{\delta_h H}{\bar{L}} \right)^{\frac{\alpha_L}{1 - \alpha_L}}$$

The price of land p_L is an exponential function of the per-unit housing price p_h which can be recovered after substituting the housing investment function (2.5) into the first order condition with respect to land, resulting in

$$(2.21) \quad p_L = \alpha_L (1 - \alpha_L)^{\frac{1 - \alpha_L}{\alpha_L}} (p_h A_h)^{\frac{1}{\alpha_L}}$$

Finally, finding the equilibrium rental rate p_r is more challenging as the market clearing condition involves a distribution, which is an infinite dimensional object. Nonetheless, it can be recovered computationally by exploiting the fact that rental demand is decreasing while rental supply is increasing in p_r , as shown below.

2.2.4. Model intuition: a supply & demand explanation

Before we turn to the two experiments performed with the help of our model economy, we present the intuition for how reducing the access to credit, either through tighter borrowing limits or via higher interest payments, affects the homeownership rate as well as house and rental prices. To that end, we use a supply and demand framework where we plot relative rather than absolute prices and quantities to capture the rent vs. owning margin, as in Greenwald and Guren (2024). Figure 2.1 plots the share of renters (demand) or the share of landlords adjusted by the number of properties they own (supply) on the x-axis against the rent-to-price ratio on the y-axis.

Rental demand, depicted by the blue line, is downward sloping because increases in the rental to house price ratio incentivize homeownership and consequently less and less households are willing to rent. On the other hand, such increases in the rent to house price ratio make buying buy-to-let properties more attractive, and as result more and more households are willing to become landlords. This results in an upward sloping rental supply curve (red line). As standard, the intersection of these two curves forms an equilibrium which determines relative price and quantities.

Now, consider the impact of a negative credit shock associated, for example, to the introduction of macro-prudential mortgage limits. On impact, the reform primarily




FIGURE 2.1. Supply and demand in the rental market

NOTE. This figure shows the main mechanisms of the model through a supply and demand illustration. The demand and supply curves are computed numerically using a suitable parameterization of the model economy.

affects potential buyers that were close to the borrowing limit before and that do not qualify for a mortgage after. These households who are not able to buy a house anymore will be pushed into renting. This shifts the demand curve outwards as shown by the blue dash line in right panel of Figure 2.1.

In a model with perfectly elastic rental supply (red dotted line), the increase in rental demand only translates into a reduction of the homeownership rate since the share of renters goes up. Prices do not move because deep-pocketed landlords are willing to buy as many houses as needed at the present value of rents to meet rental demand. However, in our model rental supply is upward sloping because the equilibrium share of landlords is endogenous. Consequently, an increase in rental demand associated with a negative credit shock results not only in a reduction of the homeownership rate, but also in an increase in the rent-to-price ratio. Moreover, this increase in the price ratio is slightly amplified in our model because landlords also use credit to buy additional rental properties, which shifts rental supply inwards (red dashed line).

Overall, a reduction in credit results in an increase of the rent-to-price ratio and a reduction of the homeownership rate because potential house buyers are credit constrained (shift in rental demand), and importantly potential landlords are sensitive to both prices (upward sloping supply curve) and credit conditions (shift in rental supply).

2.3. The Irish macro-prudential reform

Macro-prudential regulations that limit household leverage in the residential mortgage market have been widely used by policymakers to smooth the house price and credit cycles. We study the case of Ireland, whose central bank introduced these mortgage measures for the first time in February 2015 after a first discussion in October 2014. At that time, the Central Bank of Ireland established a maximum Loan-To-Income (LTI) limit of 3.5, which only applied to first-time-buyers (FTBs), and several Loan-To-Value (LTV) limits depending on the borrower and property type. For primary dwellings the limit was set to 80% of the value of the house; for FTBs, the limit was more generous: 90% for the first €220,000 and 80% for the excess amount; and for buy-to-let (BTL) properties the threshold was more stringent and set to 70%. Banks and other lenders were allowed to lend certain amounts above those limits. Specifically, for LTVs, 15% of all lending could take place above the limits, while for LTIs there was a 20% allowance. These measures have been reviewed on an annual basis since then. Nonetheless, the

alterations of these rules have been of modest nature, and the fundamental core remained unaltered until 2022. We focus on the 2015 regulation because of the prompt implementation of the reform, which, paired with data availability, makes this Irish reform a compelling case study to analyze the effects of these measures on house prices and rents in the data. We will then use this analysis as motivating evidence for the calibrated version of our model.

2.3.1. Empirical evidence

Using data for the universe of originated mortgages in Ireland, Acharya et al. (2022) study the 2015 reform and find that it generated a reduction in house price growth. In order to control for potentially confounding effects in macroeconomic aggregates, they develop a “*distance*” measure that correlates with the exposure to the macro-prudential reform. In counties where house prices were high with respect to incomes, many mortgages signed before the reform were at or above the limits: these are categorized as *low-distance* areas. One would expect that the reform would have stronger effects in these areas. In contrast, in counties where house prices were relatively low with respect to incomes, the reform was closer to non-binding and thus expected to have low to no effects. Consistently, Acharya et al. find that the “*distance*” measure positively correlates with house price growth around the reform. In other words, house prices grew more in areas where the constraints were less binding (high distance), while house price growth moderated in areas where the intervention was more binding (low distance). Non-parametric evidence of these positive correlations across Irish counties are shown in the first two panels of Figure A3.

We extend their empirical framework to analyze the effects of the reform on rental prices. We use the same “*distance*” measure and combine it with house price and rental data extracted from daft.ie (Lyons 2018). Following Acharya et al.’s empirical strategy, we regress changes in house prices and rents between the third quarter of 2014 and the last quarter of 2016 on the aforementioned “*distance*” measure. Formally, we estimate the following two regressions

$$(2.22) \quad \Delta \text{ House Prices}_i = \beta_0 + \beta_1 \text{Distance}_i + \epsilon_i$$

$$(2.23) \quad \Delta \text{ Rents}_i = \gamma_0 + \gamma_1 \text{Distance}_i + \nu_i$$

where i denotes the county, Δ is the growth rate over a 9 quarter window, and β_1 and γ_1 are the coefficients of interest. Table 2.1 shows the results of these two regressions. The

TABLE 2.1. Effect of lending limits on house and rental prices

	Δ House Prices	Δ Rents
Distance	0.289 (0.068)	-0.171 (0.039)
Obs.	54	54
R^2	0.34	0.31

NOTE. This table shows the OLS coefficients from the regressions of house price and rental price changes on the distance measure that captures the exposure to the borrowing limits.

first column replicates the positive coefficient that Acharya et al. obtain for house prices. The second column shows that the impact of the reform on rents had the opposite sign: rents increased by more in areas where the macro-prudential intervention was more binding (low distance). Quantitatively, a one standard deviation in the county-level distance measure is associated with 4.2% higher house prices and 2.5% lower rental rates. As for house prices, non-parametric evidence of these negative correlations between distance and rental price growth across counties are shown in Figure A3.




FIGURE 2.2. Placebo regression

NOTE. This figure shows the correlation between house prices and rents in response to the macro-prudential reform. Before its implementation, the ratio of OLS coefficients is positive $\beta_1/\gamma_1 > 0$, indicating co-movement between house and rental prices, while after the reform this relationship breaks and β_1/γ_1 becomes negative for a few years before it goes back to its normal positive co-movement.

To establish whether these opposite effects on rents and house prices might be the result of contrasting long-run trends in different local housing and rental markets, rather than the impact of the reform, we re-estimate equations (2.22) and (2.23) for different time windows. As in the main regression, we use the changes in house prices and rents between 9 consecutive quarters as our dependent variables while keeping the “*distance*” measure fixed to its value in 2014. We find that, consistently with our interpretation, the movement of rents and house prices in opposite directions is restricted to the time around the implementation of the policy and is not present for the rest of our sample. We report these results graphically in Figure 2.2, which shows the ratio of coefficients β_1 and γ_1 (y-axis) for the central part of each rolling window (x-axis). The unconditional average and the median value of the ratio of coefficients is positive, reflecting that, in general, house prices and rents tend to co-move. This positive co-movement is consistent with the theory as the value of a house should equal the expected discounted value of future rents. However, around the macro-prudential reform (a credit shock) this long-run relationship breaks and the ratio of coefficients becomes negative. In the short-run, potential home-buyers are constrained and pushed into the rental market which in turn generates effects going in opposite directions. The fact that the ratio of the coefficients returns to positive territory after a few years is reassuring as that the “*distance*” measure does not capture other relevant, time-invariant omitted variables, like urban vs. rural areas.

In short, this placebo test confirms that our findings are not driven by time-invariant omitted variables which are correlated with the “*distance*” measure and reinforces the idea that the credit shock induced a decoupling of the usually positive relationship between the evolution of house prices and rents.

2.3.2. Model calibration

We parametrize our model economy to be consistent with the cross sectional features of the Irish economy and use the calibrated version of the model to: (i) understand the opposite impact of the macro-prudential reform on house prices and rents and (ii) further analyze the distributional effects and the costs imposed on households by these reforms, while taking into account a broad life-cycle perspective.

As standard in the macroeconomic literature, we assign some of these parameters externally, while others are chosen internally with the objective of minimizing the distance between a collection of data and model moments.

2.3.2.1. *Externally calibrated parameters*

Demographics and Preferences. The model period is one year. Households enter the economy at age 25, they retire with certainty at age 65 and live until age 95. This means that $J^{ret} = 41$ and $J = 71$. There is no population growth. We set the CRRA risk aversion coefficient γ to 2, a common value in the literature. The scale parameter of the taste shock σ_ϵ is within the range suggested by Iskhakov et al. (2017) and equal to 0.05.

Earnings process. Our measure of income in the data is disposable household income after both taxes and transfers. We estimate our earnings process following De Nardi, Fella, and Paz-Pardo (2020). Namely, we extract out the persistent and the transitory component of earnings using the procedure described in Arellano, Blundell, and Bonhomme (2017), and then incorporate the dynamics for the persistent component in a nonparametric way. Applying this procedure allows us to estimate earnings dynamics under flexible assumptions, and in particular incorporating potential age-dependence, non-normalities and non-linearities in earnings dynamics. The first element is of particular relevance for our question. Most households become homeowners when they are relatively young, still changing jobs and potentially subject to large fluctuations to their labour market income. A standard earnings process in which earnings are a random walk is a poor representation of the earnings risk faced by households at this particular age. Middle-aged households with stable jobs, instead, have much higher persistence, but significant negative skewness risk (e.g., through job loss). For a detailed description of the method and the economic implications of flexible earnings dynamics, see De Nardi, Fella, and Paz-Pardo (2020).

We use data from the Household Finance and Consumption Survey (HFCS) to extract the average age-earnings profile in the Irish economy after taking into account year effects. The HFCS, which takes place every three years, collects rich data on the income and wealth of European households, including their homeownership status, rental income, etc., which we also use as targets for our model. However, the triennial nature of the survey does not allow us to estimate an annual earnings process. Hence, for the stochastic component of the earnings process we use household earnings data for the United Kingdom from De Nardi, Fella, and Paz-Pardo (2024), who extract them from the BHPS/Understanding Society survey, and assume that the stochastic properties of household earnings are similar across both countries.³

³ We have also estimated an annual earnings process for Ireland based on EU-SILC data (European Union Statistics on Income and Living Conditions), which, despite being nationally representative, is

Housing. We set the maximum amount of houses that a household can own X to 3 both for simplicity and computational considerations. However, this choice implies that we can capture the vast majority of landlords in Ireland: over 80% of them have only 1 or 2 rental properties, and they represent over half of all tenancies. Indeed, around 37% of all rental properties are owned by households with just one buy-to-let property – see Figure A2 – and over 60% of non-occupier transactions in 2015 involve household buyers – see Figure A1. Large institutional investors (non-household buyers) have grown in relevance over the past decade, but they were still relatively small players when the reform took place.⁴ Since then, institutional investors have been mostly concentrated in a set of particular rental submarkets (e.g. new builds in highly sought-after areas). Given that we do not model this spatial heterogeneity, we abstract from them in our framework.

We assume that there are two house qualities and normalize the lowest quality $\tilde{h}_1$ such that the aggregate amount of housing H is equal to 1 in equilibrium. The value for the better quality $\tilde{h}_2$ is chosen such that the ratio of prices is fixed and equal to the owner-occupied to buy-to-let price ratio in the data. These two different type of houses not only differ in their final transaction price, but also in the housing utility flow they report to those living in them. We assume that households get a premium from ownership which is increasing and concave in the quality of the house with $\alpha_h = 0.5$ controlling the curvature of this function.

The housing depreciation rate δ_h is set to be 1.2% per year and it is within the range of typical values used in the literature. The transaction cost for selling and buying a house τ_h equals 3% of its value. The maximum loan-to-value, λ_{LTV} , and loan-to-income, λ_{LTI} , ratios before the macroprudential reform are 1.0 and 6.0, respectively. This is consistent with the evidence in Kelly, McCann, and O’Toole (2018) that estimate the 98th percentile of observed LTI and LTVs ratios on quarterly mortgage data during the period 2003 to 2011. Note that prior to the 2015 reform, there were no institutional limits. Therefore, these limits correspond to those that were imposed by Irish banks based on their own risk assessment.

targeted to produce statistics on poverty and living conditions and hence might capture the earnings dynamics of the upper part of the income distribution in a more limited way. Our main results with this alternative earnings process are very similar.

⁴ Oosthuizen (2023) explains this upward trend in the US by the decreasing operating costs of larger institutional investors, while Muñoz and Smets (2022) argue that it is associated to the lack of regulatory lending limits to real estate funds.

TABLE 2.2. Parameter values

Parameter	Interpretation	Value
<i>Externally calibrated</i>		
J^{ret}	Working life (years)	41
J	Length of life (years)	71
γ	Risk aversion coefficient	2.0
σ_ε	Taste shock scale parameter	0.05
X	Maximum amount of houses owned	3
$\{\tilde{h}_1, \tilde{h}_2\}$	Housing qualities	{0.905, 1.095}
α_h	Curvature in utility premium function $f(\cdot)$	0.5
δ_h	Housing depreciation rate	0.012
τ_h	Proportional transaction cost	0.03
λ_{LTV}	Maximum loan-to-value ratio	1.0
λ_{LTI}	Maximum loan-to-income ratio	6.0
r	Risk-free rate	0.02
κ	Intermediation wedge	1
A_c	Aggregate labor productivity	1.2055
$\bar{L}$	Amount of land	1.0
α_L	Share of land in production	0.33
ξ	Adjustment cost scale in housing production	0.16
<i>Internally calibrated</i>		
β	Discount factor	0.9656
θ	Utility premium from living in a low quality house	1.3378
A_h	Scaling factor in housing production	0.121

NOTE. This table shows the value of the parameters used for solving our model economy and to carry out the experiments. For the macro-prudential intervention, λ_{LTV} and λ_{LTI} will change; while for the increase in rates, r and κ will be adjusted.

Financial instruments. The risk-free rate on liquid savings r_s is set to 2% per annum. The proportional wedge κ is set to 1, implying a mortgage rate r_b of 4% per annum. This is consistent with the gap between the average mortgage rate and the 10-year yield on government debt.

Production. The final good aggregate productivity shifter A_c is set to 1.2055, which is also the average wage in the economy. The amount of buildable land $\bar{L}$ is normalized to 1, and the share of land used in production in the housing sector α_L is fixed to 0.33, which implies an elasticity of housing supply of approximately 3.

2.3.2.2. Internally calibrated parameters, targets, and model fit

The remaining three parameters: the discount factor β , the homeownership utility premium for living in a small/low quality house θ , and the scaling factor in housing production A_h , are jointly chosen to match four moments of the data. We target the average wealth to income ratio, which is around 7 in the HFCS; the homeownership rate that was on average around 80% according to EU-SILC; a house price to income ratio of 5 that is consistent with the data in the Central Statistics Office (CSO); and the house price to rent ratio that is computed using data from the Residence Tenancies Board (RTB) and the CSO. The first block of Table 3.2 shows the exact value of these four moments in the data as well as their model counterparts, which were obtained using the parameters in the last block of Table 3.1.

The model is able to match the average homeownership rate, the average house price to income ratio and the house price to rent ratio reasonably well. However, it slightly under-predicts the average wealth to income ratio. More importantly, the model is able to replicate the share of landlords in the economy, both at the aggregate level and along the age distribution. At the aggregate, the share of landlords with two rented out properties (the upper bound) is 4.30% in the model; while in the data, 5.11% of landlords own two or more rental properties. Along the life-cycle, it is only at mid age when a significant fraction of households can afford to buy a second or third home, consistently with the data. Later in life, households in the model, unlike those in real

TABLE 2.3. Targets and model fit

Moment	Model	Data	Source
<i>Targeted</i>			
Wealth to income ratio	5.89	6.78	HFCS
Homeownership rate	79.42%	80%	EU-SILC
Avg. house price to income ratio	4.93	5.0	CSO
House price to rents ratio	22.73	22.58	RTB/CSO
<i>Untargeted</i>			
Rents to avg. income ratio	0.196	0.2216	RTB/CSO
Share of households with 3+ properties	4.30%	5.11%	HFCS

NOTE. This table shows the model ability to capture certain features of the Irish economy. The top block corresponds to the targets used in a minimum distance estimation, while the bottom block show the performance of the model relative to untargeted moments.




FIGURE 2.3. Number of properties along the life-cycle

NOTE. This figure shows the number of properties owned by households in the data (panel A) and in the model (panel B) at different ages. Since the data is aggregated in 5 year groups, we also aggregate it in the model for ease of comparability.

life, sell these properties to finance retirement – see Figure 2.3. This mismatch is a standard feature of life-cycle models which do not include a set of relevant features of retiree saving behavior, such as precautionary savings related to medical costs or long-term care and bequest motives (Nakajima and Telyukova 2020). As also shown in Figure 2.3, the share of renters decreases with age in both the model and in the data. However, this happens more quickly in the model than in real life. In any case, these life-cycle patterns are endogenously captured by the model without explicitly targeting them, which is reassuring about the validity of the model as a laboratory to study the distributional effects of the macro-prudential reform discussed above.

Rental supply elasticity. An additional moment that is informative about the ability of the model to generate realistic changes in house prices and rents as a response to policies and shocks is the elasticity of rental supply, i.e. by how much rental prices need to increase in order to encourage landlords to supply 1% more units of rental housing, which is closely related to the slope of the rental supply curve that we represent as a solid red line in Figure 2.1. In our model, this elasticity is an endogenous object that depends on the distribution of wealth and income of landlords in the economy and their policy functions.

Identifying rental supply elasticities in the data requires, ideally, a shock that affects rental demand alone without having any impact on house prices or rental supply. As a result, there are no readily available estimates for the Irish case that we can use as indicators of model fit for our experiment. Rotberg and Steinberg (2024) estimate this parameter for the United States using information on the incidence of property taxes, and find that it is 1.4 in the long run (a 1% increase in rental prices leads to a 1.4% increase in quantity supplied), with some other empirical studies suggesting that it might be even lower. In comparison, they show that this elasticity is infinite in models with rental sectors in which prices are determined by a user-cost formula (as price to rent ratios are independent of quantities), and that it is very large in models in which landlords are homeowners that choose every period how much of their housing stock to rent out (e.g. 38 in Floetotto, Kirker, and Stroebel (2016)).

The elasticity of rental supply in our model is 3.5 in our pre-reform steady state – larger than that in Rotberg and Steinberg (2024), but an order of magnitude smaller than in other models with endogenous landlords. The reason is that, in our model, housing is lumpy and illiquid also for landlords, who need to buy one complete house to be able to rent it out. Thus, there are frictions associated with becoming a landlord, which makes them less responsive to changes in rent-to-price ratios than those in Floetotto, Kirker, and Stroebel (2016) or Sommer and Sullivan (2018). Additionally, and compared to the rental company in Rotberg and Steinberg (2024), our model generates a nonlinear rental supply function with elasticities varying along the curve, derived from the fact that the wealth, income and age of the marginal landlord change as we move along the supply curve.

2.3.3. Aggregate and distributional effects of tighter borrowing limits

We study the effects of the macroprudential reform under the assumption that it is a permanent change. We begin by comparing two steady state equilibria that only differ in their mortgage borrowing limits.⁵ Then, we consider the effects of the transition from the (initial) *pre-reform* steady state to the (final) *post-reform* steady state. That is, agents unexpectedly observe that borrowing limits become more stringent but after this first initial surprise they are aware of such permanent change in credit conditions. Finally, we use these results to evaluate the welfare effects of the reform on our heterogeneous households.

⁵ Appendix B.3.2 shows the effects of tightening LTV and LTI limits in isolation. Given our parameterization, the LTI limit has a stronger effect than the LTV but there are interactions between the two.

2.3.3.1. *Steady state comparison*

In the *pre-reform* economy, households are able to borrow up to 100% of the value of their house and up to 6 times their annual income. Equilibrium quantities and prices under these credit conditions are reproduced in the first column of Table 2.4. The second column presents the equilibrium outcomes in an economy where these limits correspond to the institutional ones introduced by the 2015 reform. Hence, in the *post-reform* economy, households face: (i) a 80% loan-to-value limit if they buy an owner-occupied property, (ii) a 70% loan-to-value limit if they purchase a buy-to-let home, and (iii) a 3.5 loan-to-income limit if they are first time buyers. In the model we identify owner-occupied and first time buyers as those households that move from being renters into homeowners. Similarly, buy-to-let purchases correspond to those carried out by households that transition from homeowner to landlord as well as those from landlords expanding their real estate portfolio. All other borrowers face the borrowing constraints that were in place in the *pre-reform* economy and capture limits imposed by banks based on their own risk assessment.

As a result of these changes, and consistently with the intuition in Section 2.2.4, the homeownership rate falls by 1.83 percentage points while the rent-to-price ratio increases by 2.82% in the long-run. Note that such increase in the rent-to-price ratio could be explained by either an increase in rents or a fall in house prices. Importantly, our model is able to disentangle these two effects. We find that most of this increase in the price ratio is driven by the rental price as it *increases* by 2.74%, while the average house price in the economy stays more or less *constant* and falls only by less than 0.01%.

These price dynamics arise because: (i) the marginal landlord needs to be compensated even more, via higher rental prices, to meet the increased rental demand coming from households that do not qualify for a mortgage at the new borrowing limits, and (ii) some households are forced to buy lower quality houses (downsize). As this second effect is rather limited – there is only a tiny increase in the share of low quality homes (0.38 p.p.) – the reduction in the average house price coming from a composition effect with a higher share of cheaper low quality homes is very small. Moreover, we find that the marginal landlord is typically a homeowner that becomes landlord for the first time rather than already existing landlord that expand their real estate business, as more than 75% of the new rental demand is covered by households that weren't landlords in the pre-reform economy ($0.22 \times 2/1.83 \approx 0.25$). Figure A4 summarizes these flows across different housing tenure statuses.

TABLE 2.4. A credit crunch – tightening LTI & LTV limits

	Pre-Reform	Post-Reform	Percentage Change
Rent-to-Price	3.98 %	4.09 %	2.82 %
Average house price to income	4.930	4.925	-0.01 %
Rent to Income	0.196	0.201	2.73 %
Homeownership rate	79.42 %	77.59 %	-1.83 p.p.
Share of households living in high-quality homes	50.41 %	50.03 %	-0.38 p.p.
Share of households with 3 properties	4.29 %	4.51%	0.22 p.p.

NOTE. This table show the equilibrium prices and quantities for the two economies considered: (i) the pre-reform economy in which borrowing limits are loose ($\lambda_{LTV} = 1$ and $\lambda_{LTI} = 6$), and (ii) the post-reform economy in which tighter borrowing limits are imposed ($\lambda_{LTV}^{oo} = 0.8$, $\lambda_{LTV}^{btl} = 0.7$ and $\lambda_{LTI} = 3.5$).

As a result of the reform, the concentration of housing wealth rises in two ways. First, the increase in the number of renters implies that their homes are now owned by landlords, who are already homeowners themselves. Second, the number of landlords that hold two rental properties, which in our framework proxies relatively larger owners, increases by 0.22 percentage points as a result of the reform.

2.3.3.2. *Transition dynamics & welfare*

The steady state comparison is extended to consider the transition from the *pre-reform* to the *post-reform* economy. This analysis will allow us to determine which household groups benefited and which ones lost from the introduction of tighter LTV and LTI limits.

Transition paths. Figure 2.4 illustrates the path of rental and house prices as well as the evolution of the homeownership rate during the transition to the steady state in which borrowing limits are set to those imposed by the Central Bank of Ireland in 2015. Following the macro-prudential reform, prospective homeowners cancel or postpone their buying decisions and stay renters, increasing the demand for rental accommodation. As a result, rental prices initially jump near the new steady state equilibrium price to incentivize the (endogenous) landlord formation that meets the excess demand. The homeownership rate declines, but does not immediately adjust to its new steady state level as new born generations face a more difficult environment to get on the property ladder and stay in the rental market for longer. Consequently, the rental price has to further increase and reaches its maximum level after 4 years – a 3.61




FIGURE 2.4. Transition paths: from loose to tight credit limits

NOTE. This figure shows the evolution of the rental price (top-left), the average house price (top-right) and the homeownership rate (bottom) along the transition from the old to the new steady state.

% increase relative to the pre-reform – before slowly going back to its new steady state level. Average house prices also fall initially as the composition of the housing stock changes and features a slightly higher fraction of cheaper low quality homes. Given that this compositional effect is stronger than the change in the pre-unit house price, which jumps below the post-reform level and reverts back to its new steady state level quickly, the average house price doesn't overshoot and slowly converges to the new steady state level. These price paths are consistent with the negative empirical correlation of tighter macro-prudential borrowing limits and house prices, as well as the positive empirical correlation of these limits and rental prices.

Welfare. We evaluate the distributional effects of the reform through the traditional lifetime consumption equivalent variation (CEV) measure. This metric informs us about how much consumption (in percentage) needs to change in the pre-reform economy such that the households are indifferent between living in the pre-reform steady state and living through the transition induced by the policy reform. Formally, for a given set of state variables $x = (a, y, h, j)$, we compute the consumption equivalent variation $g(x)$ as

$$(2.24) \quad V_0(x; g) \equiv (1 + g)^{1-\gamma} V_0(x) = V_1(x) \quad \Rightarrow \quad g(x) = \left[\frac{V_0(x)}{V_1(x)} \right]^{\frac{1}{1-\gamma}} - 1$$

where we are using the fact that the utility function is CRRA. From (2.24) it is easy to realize that a negative value of $g(x)$ is associated with agents being worse-off by the introduction of the reform. Figure 2.5 depicts the value of this metric along the income distribution (panel A) and household's age (panel B).

Because welfare is affected by the tighter limits as well as by the associated price movements, we decompose the CEV into partial and general equilibrium effects. The dashed line depicts the welfare effects of the macro-prudential reform in absence of price movements. As expected, the tightening of borrowing limits in itself is welfare reducing for all agents in this economy as we are constraining the feasible set. Young and middle-income households bear most of the costs given that they are prospective homeowners.

Turning now to the price effects, rental prices need to adjust upwards to incentivize more household to become landlords and cover the additional rental demand. As a result, there is a further welfare loss experienced by young, poor and middle income households. These household are forced to pay higher rental prices reducing their savings and cash available for consumption. Graphically, this effect is shown by the




FIGURE 2.5. Price adjustments & welfare

NOTE. This figure plots the value of the CEV along the income (panel A) and age distributions (panel B) and decomposes the overall welfare effect on the contribution of the reform itself (tighter borrowing limits) and the price adjustments (higher rents, lower house prices).

gap between the dashed and the dotted line, which is shrinking along the income distribution and it is even positive, and hence welfare improving, for the very rich. In fact, those at the top of the income distribution as well as middle age households benefit from the increase in rents because they are typically landlords and hence receive a larger cash flow from their real estate investment. Finally, the welfare effect of house prices is rather limited as they remain nearly constant during the transition. Graphically, this is shown by the overlap between the solid and the dotted lines.

Finally, we decompose the CEV based on households's housing tenure. As we have just seen, the macro-prudential reform has a direct impact on prospective homeowners and it also indirectly affects renters and landlords through price adjustments. In fact, housing tenure status is a great indicator to disentangle the winners and losers of constraining credit via more stringent LTV and LTI limits, as shown in Figure 2.6. In a nutshell, renters lose, homeowners are indifferent and landlord benefit. Consequently, a welfare neutral policy would require a redistribution of consumption from landlords to renters to compensate for the unintended effects of the macro-prudential reform.




FIGURE 2.6. Housing tenure & welfare

NOTE. This figure plots the value of the CEV along the age distribution for three group of households: renters, homeowners and landlords.

2.4. Interest rates, credit standards, and price dynamics

In this section, we use the model presented in Section 2.2 to study the effects of a different type of credit shock: an exogenous, unexpected and permanent rise to the real interest rate of 1 percentage point. We do so under the assumption that the economy has the macro-prudential measures described in Section 2.3 in place, and then show how effects would differ under a looser credit conditions.

2.4.1. A permanent rise in the real interest rate

An increase in the interest rate has a direct impact on prospective homeowners as it makes mortgage credit more expensive. In this regard, this experiment is similar to the tightening of LTV and LTI limits studied in Section 2.3. Unlike the macro-prudential reform, the increase in interest rates also has a direct impact on savers because the return on bonds increases, and on current mortgagors because their mortgage interest payments also increase.

In Table 2.5, we show the joint impact of all of these channels on the rent and price to income ratios, the homeownership rate, and the share of high quality homes. By comparing the *low interest rate* and *high interest rate* economies shown in the leftmost and rightmost columns, we observe that the increase in the real interest rate reduces the homeownership rate (-0.92 p.p.) and increases the rental rates (12.7%). Unlike in our previous experiment, the drop in average house prices is significant (-1.62%) and many households choose to downsize and buy smaller houses. Similar to Figure A4, Figure A5 shows the flows across different housing tenure status for the interest rate experiment and in particular the larger movement from big to small owners relative to the macro-prudential credit tightening.

In order to understand and decompose the different channels that drive these results, we proceed in steps and study the effect of the increase in the savings rate and the mortgage rate in isolation. To do so, we introduce a counterfactual economy in which the return on savings is $r_s = 3\%$ but the mortgage rate is still $r_b = 4\%$.

2.4.1.1. An increase in the return on savings versus a rise in the borrowing rate

A permanent rise in the return on savings. To analyze the effect of an increase in the rate of return of savings alone, we compare the *low interest rate* economy (1st column, Table 2.5) and the counterfactual economy (2nd column, Table 2.5). Under tight credit

TABLE 2.5. Increasing the real interest rate

	Low Int. Rate	Decomposition	High Int. Rate
	$r^s = 0.02, r^b = 0.04$	$r^s = 0.03, r^b = 0.04$	$r^s = 0.03, r^b = 0.05$
<i>Tight credit conditions</i>			
Rent-to-Price	4.09 %	4.58 %	4.69 %
Average house price to income	4.925	4.899	4.846
Rent to Income	0.201	0.224	0.227
Homeownership rate	77.59 %	76.99 %	76.67 %
Share of households living in big houses	50.03 %	47.74 %	43.02 %
<i>Loose credit conditions</i>			
Rent-to-Price	3.98 %	4.48 %	4.57 %
Average house price to income	4.930	4.880	4.835
Rent to Income	0.196	0.219	0.221
Homeownership rate	79.42 %	78.93 %	78.35 %
Share of households living in big houses	50.41 %	46.01 %	42.02 %

NOTE. This table show the equilibrium quantities and prices for three economies: (i) low interest rate, (ii) high interest rate and (iii) a counterfactual economy that decomposes the effects of the return on savings and the borrowing rate. Each of these economies are analyze under different credit conditions: tight credit (top block) and loose credit (bottom block).

conditions, rising the saving rate leads to a 0.6 p.p. fall in the homeownership rate. Financial assets become relatively more profitable than housing, *ceteris paribus*, thus generating a substitution effect that decreases the incentives to own a house. This effect is stronger than the higher returns from savings that allow prospective homeowners to save for downpayment at a much faster rate (income effect). The substitution effect is particularly strong for landlords, who do not get utility from any additional properties and thus treat them as a pure financial investment. As a result, they require a very large rise in the rental rate, higher than 11%, to meet the increased rental demand. This channel explains the bulk of the overall increase in rental prices that we observe as a result of the permanent rise in r . Moreover, the negative substitution effect between housing and financial assets also results in many households buying a smaller house, which puts downward pressure on aggregate house prices (-0.5%).

A permanent rise of the mortgage rate. To study the effects of an increase in the mortgage rate alone, we compare the counterfactual economy (2nd column, Table 2.5) to the *high interest rate* economy (3rd column, Table 2.5). The rise in the borrowing rate has a negative income effect on housing demand, as it increases the interest payments of

current mortgagors, but also a negative substitution effect, as it increases the cost of accessing credit. As a result of the rise in borrowing costs, the homeownership rate falls by 0.33 percentage points and there is a sizable mass of households that opt for the less expensive house – the share of low quality homes increases by 4.72 percentage points. Consequently, the average house price in the economy falls by 1.1%, which is more than twice as large as the fall associated to the increase in the return on savings. On the other hand, the increase in the rental price associated with the increase rental demand is of an order of magnitude smaller and it only rises by 1.2%.

In our model, we assume that all mortgagors have a floating-rate mortgage and thus the adjustment to their mortgage payments happens immediately after the shock. This assumption is consistent with the Irish institutional structure, in which around 80% of mortgages were de facto adjustable rate in 2018 (Badarinza, Campbell, and Ramadorai 2018), with some of them being fully adjustable and some of them having fixed rates for a relatively short period, such as 1 or 2 years. In an economy where fixed-rate mortgages are more prevalent, such as the United States, the steady state impact would be similar, but over the transition existing mortgagors would be less negatively impacted.

2.4.2. Interaction with credit standards

In the results we have shown so far, we have analyzed the effects of the interest rate increase in an environment in which institutional limits on mortgage borrowing were in place. To gauge to which extent macroprudential policies can help cushion the aggregate effects of shocks such as an unexpected interest rate rise, we also study the effects of the same shock in a context in which borrowing limits are less stringent and equal to those prevailing in Ireland before 2015. The results from this last experiment are shown in the bottom panel of Table 2.5.

We find that the fall in the homeownership rate (-1.07 p.p.) as well as the fall in average house prices (-1.93%) are larger under looser credit conditions. On the other hand, the rise in the rental price is of similar magnitude as it rises by 12.84% under looser credit conditions. Prospective landlords are less affected by the borrowing limits than prospective homeowners, and consequently the adjustments in the housing market are more sensitive to the credit conditions than those in the rental market. These results suggest that macroprudential policies, despite the negative welfare effects we have discussed in Section 2.3, can be effective at mitigating the fluctuations in housing markets originated by other types of shocks.

2.4.3. Transition dynamics & welfare

We now turn to studying how the change in the real interest rate impacts rents and house prices over the transition, and then look into the welfare of the households who are impacted by the shock. For this section, we assume that tight LTI and LTV limits are in place already in the initial steady state.

Transition paths. Figure 2.7 depicts the equilibrium paths for the rental price, the average house price and the homeownership rate after an exogenous permanent change in the real interest rate. Similarly to our previous experiment, rental prices overshoot and reach its peak after 4 years with an increase of 20.8% relative to the low interest rate economy. Nonetheless, the convergence to its new steady state level is slower and takes about 20 years. On the other hand, the average house price adjustment is relatively faster, but still takes about 15 years to reach its new lower level. The homeownership rate jumps to its new level within the first couple of years and stays there for the remaining of the transition.




FIGURE 2.7. Transition paths: from low to high real interest rates

NOTE. This figure shows the evolution of the rental price (top-left), the average house price (top-right) and the homeownership rate (bottom) along the transition from the low to the high interest rate economy.

Consumption Equivalent Variation. Figure 2.8 plots the welfare impact for households who are initially young (25-30 years old) when the interest rate shock takes place, decomposing it into interest rate changes (borrowing rate vs. return on savings) as well as into house and rental price adjustments. The overall impact, depicted by the solid line, shows that there are winners and losers from the permanent increase in the real interest rate. In particular, those below the median of the income distribution are worse off, while those above are better off. Adding one channel at a time, we observe first that all households benefit from the rise of the return on savings (dash-dotted line), but high income households benefit comparatively more because they have more wealth. Adding the increase in the cost of mortgages (dashed line) reduces the welfare gains for some parts of the distribution, particularly those in middle income deciles who need to acquire larger mortgages to buy a house, but less for the relatively richer, who are more likely to buy homes outright or with smaller mortgages, and for the relatively poorer, who are mostly renters. Like in our previous experiment, the general equilibrium effect coming from price adjustments in the rental markets is relatively large and has a very negative welfare impact on low income households, who as a result lose after the shock (dotted line). Although the drop in house prices is larger in this experiment, its welfare effects are relatively small in comparison.



FIGURE 2.8. Rate changes, price adjustments & welfare along the income distribution

NOTE. This figure shows the consumption equivalent variation for a group of young households along the income distribution. The dashed-dotted line shows the effect of increasing the return on savings alone, the dashed line factors in the increase in the borrowing rate, while the dotted line takes into account the adjustment of the rental price. Finally, the solid line shows the overall effect of the increase in the real rate in terms of welfare.

2.5. Conclusion

In this paper, we build an equilibrium model of the housing and rental markets in which households differ in their age, income and wealth as well as in their housing tenure status (renters, homeowners, or landlords). Endogenous landlord formation allows us to have an upward sloping rental supply curve that is crucial to understand the effect of credit conditions on house and rental prices. We use this model to analyze two shocks to the availability of credit for households: (i) the introduction of LTI and LTV limits in Ireland in 2015, and (ii) a permanent rise in the real interest rate.

Regarding the 2015 Irish macro-prudential reform, we show empirically that it had opposite effects on house prices and rents. In the model, these effects are explained by a rise in rental demand that needs to be met by a landlord sector which is heterogeneous and constrained, and that as a result displays an upward sloping rental supply curve. Quantitatively, most of the adjustment occurs through the increase in rental prices (2.74% upon impact and 3.61% after 4 years) with house prices not reacting much to the tightening of LTI and LTV limits in the long run. Nonetheless, adjustments via quantities in the housing market were sizable as the homeownership rate fell by 1.8 p.p in the long run. These changes impacted households welfare differently. Renters, who are typically young, poor or middle-income households, suffered the most because they had to pay higher rents and were forced to postpone or cancel their buying decisions. On the other hand, landlords who are top income earners benefited as they are not constrained by the new limits, they can get slightly cheaper houses and receive higher cash flows from their real estate businesses.

Turning to our second experiment: the rise in the real interest rate, we find that it also leads to a reduction of the homeownership rate and the average house price as well as to an increase in rental prices. Moreover, these price adjustments in the housing and rental markets generate losers and winners, highlighting the potential redistributive effects of interest rate shocks or monetary policy more broadly via the housing and rental markets. In particular, middle-income and top earners benefit from the higher return on savings as well as the lower house prices, while poor households lose from the higher rental prices.

Our paper highlights that rental markets are key to understand the equilibrium impacts and welfare effects of credit shocks, and provides a theoretical and quantitative framework to analyze shocks in which both house prices and rents can react endogenously and in potentially different directions. These results open interesting

avenues for future research. For instance, although in our second experiment we focus on an exogenous real interest rate shock, it is likely that similar channels operate as a reaction to monetary policy shocks in a model with nominal rigidities, which could help us to better understand its transmission channel through the housing market. Besides, our model assumes that there is one single housing and rental market at the national level, with different housing qualities but a common price. Consequently, an interesting extension would consider spatial heterogeneity across different urban and rural areas, and the interaction of a credit shock in a context of accelerated urbanization and increased house price and rent inequalities.

Chapter 3

The Role of Mortgage Interest Fixation Periods for Monetary & Macro-prudential Policies

Joint with Stephen Millard & Alexandra Varadi

Abstract: In many countries the most common mortgage contract neither has a fixed nor a fully adjustable rate. The typical interest fixation period varies between two to ten years. This paper offers a structural analysis of how the interest rate fixation period affects the transmission of monetary policy and how it interacts with borrower-based macro-prudential limits. Using a general equilibrium model with long-term mortgage debt and calibrated to the United Kingdom, we find that: (i) the interest fixation period and the tightness of credit conditions do not matter if monetary policy shocks are transitory, (ii) looser credit limits and shorter fixation periods amplify the redistributive effects of inflation target shocks that increase nominal rates persistently, and (iii) LTV limits act as a backstop to the high sensitivity of PTI limits to monetary policy, especially when the interest fixation period is short.

3.1. Introduction

The relevance of household's balance sheets for the transmission of monetary policy dates back to the net worth channel of Bernanke and Gertler (1995). After the Great Recession, it has gained even more popularity as housing was not only the main driver of the boom and bust, but also was at the center of the transmission of monetary policy into household's consumption. Note that housing is the largest asset in most households' portfolios, and importantly, it is also typically associated to a mortgage as its value often exceeds households' net worth (Campbell and Cocco 2003, Piazzesi and Schneider 2016). Moreover, it is well-known that outright owners' consumption is affected by changes in house prices, either via the housing wealth channel (Iacoviello 2011, Aladangady 2017) or the housing collateral channel (Cloyne, Huber, Ilzetzki, and Kleven 2019); while mortgagors' consumption is in addition affected by changes in their mortgage interest payments (Di Maggio, Kermani, Keys, Piskorski, Ramcharan, Seru, and Yao 2017). As a result, housing tenure status is a key determinant of the strength of the transmission mechanism from interest rates into consumption (Cloyne, Ferreira, and Surico 2020). The strength of these effects also depend substantially on the interest rate schedule associated to the life cycle of the mortgage loan as well as on the credit limits imposed on newly issued mortgage debt. After the Great Recession, policymakers also started to design and implement borrower-based macro-prudential policies to prevent or smooth the impact of future shocks through the housing market, however, little is known about how these limits interact with the persistence of monetary policy shocks as well as with the interest rate schedule of the typical mortgage.

In this paper, we try to fill this gap and offer a structural equilibrium approach to answer the following questions: *how does the strength of monetary policy depend on the mortgage interest fixation period? And how it is affected by credit conditions?* After documenting that in the United Kingdom the typical interest fixation period is either two or five years and that similar fixation periods also predominate in many other countries, we build a general equilibrium model based on Greenwald (2018) in which we compare the transmission of monetary policy into consumption under three different mortgage contracts: fixed rate mortgages (FRM), adjustable rate mortgages (ARM) and hybrid rate mortgages (HRM). The latter are the main theoretical innovation in this paper and allow mortgage payments to switch from fixed to adjustable rates at different times in the life-cycle of the mortgage. Similarly to Garriga, Kydland, and Šustek (2021), we use these counterfactual economies to understand how temporary and persistent




FIGURE 3.1. Evolution of Bank & Mortgage Rates

NOTE. Panel A of this figure shows the policy rate chosen by the Bank of England during the period from 2010 to 2023. On the other hand, panel B shows the evolution of mortgage rates for various interest fixation periods and a 75% loan-to-value (LTV) limit.

monetary policy shocks propagate into the real economy. Differently from their paper, we also test the sensitivity of monetary policy transmission to housing tools. First, we investigate how having Payment-to-Income (PTI) limits on top of Loan-to-Value (LTV) limits may dampen or amplify the responses to these shocks. Note that PTI limits are directly affected by changes in interest rates, while LTV limits only change indirectly through the effect of interest rates on house prices. And second, we test how different LTV and PTI calibrations may impact the effects of monetary policy on house prices and the real economy. Moreover, we also focus on interest rate hikes rather than drops as most Central Banks (CBs) have raised their interest rates in the past couple of years. For example, the Bank of England (BoE) raised their interest rate from 0.1% to 5.25% between December 2021 and August 2023, which in turn translated into an increase of 300 basis points in mortgage rates – see Figure 3.1.

Our findings can be summarized in two sets of results. First, relative to the length of the interest fixation period and its impact on the monetary policy transmission, we find that: (i) the response of consumption, output and inflation to a *transitory monetary policy shock* is independent of the mortgage interest fixation period as the standard New-Keynesian channel dominates, leading to increased savings and reduced investment after an increase in the risk-free rate; and (ii) a shock to the inflation target that leads to *persistently higher nominal interest rates*, but otherwise, has no significant changes in the real economy, has important redistributive effects from borrowers to savers which are stronger in the *ARM economy* as mortgage rates react one-to-one with the policy

rate, leading to a larger increase in mortgage payments. For the *FRM economy* such an increase is gradual, while it is delayed until the end of the interest fixation period for the *HRM economy*. Consequently, it is less costly for the borrower in these two economies to insure against the higher mortgage payments.

The second set of results speak about the interaction with macro-prudential tools. In particular, we find that: (i) a shock that leads to persistently higher nominal interest rates is more powerful under looser credit conditions, and in particular, under loose LTVs, which amplify the redistributive effects of the shock, independently of the interest fixation period; and (ii) the tightness of credit limits is irrelevant when monetary policy shocks are transitory. These two findings do not imply that we should not care about the interest fixation period when analyzing the interaction between monetary and macro-prudential policies as these results rely heavily on the split between LTV- and PTI-constrained borrowers. In fact, in a *PTI only economy*, the response of mortgage debt issuance to a temporary shock is twice as large in the *ARM economy* relative to the *FRM economy*. House prices also fall approximately 2 percentage points more. On the other hand, these differences are not present in a *LTV only economy*. In other words, the presence of LTV limits smooths out the effects associated with the very sensitive PTI limits to interest rate changes to the extent that the different pass-through associated to different interest fixation periods is deemed to be irrelevant.

Related Literature. This paper is closely related to theoretical papers that analyze the role of household mortgage debt in the transmission of monetary policy, e.g. see Iacoviello (2005), Garriga, Kydland, and Šustek (2017), Greenwald (2018), Auclert (2019), Wong (2019), Beraja, Fuster, Hurst, and Vavra (2019), Garriga, Kydland, and Šustek (2021), Berger, Milbradt, Tourre, and Vavra (2021), or Eichenbaum, Rebelo, and Wong (2022). In several such papers, there is a comparison between ARMs and FRMs given the different pass-through of the policy rate. Recall that in ARM contracts the mortgage rate is linked to the short term nominal interest rate, while in FRM contracts there is a constant rate set at origination. As a result, changes in the policy rate affect mortgage payments almost immediately under ARMs, while it only has an effect on newly issued loans under FRMs. Garriga, Kydland, and Šustek (2021) explore this distinction through the lens of a New Keynesian model with long-term mortgage debt. They show that if monetary policy shocks are transitory, then firms' output responses dominate other channels and consequently the mortgage contract plays a minor role in shaping the responses of macroeconomic aggregates. On the other hand, if monetary shocks are

persistent, then firms react through prices rather than output, which in turn affects real mortgage payments and generates redistributive effects, especially under ARMs. Our paper complements their work by analyzing other interest fixation periods beyond the two extremes, FRM and ARM, as well as by considering the interaction with housing tools, which we show that matter for the transmission of persistent shocks. Greenwald (2018) also studies the responses of aggregate variables to transitory and near-permanent shocks to mortgage rates through a similar general equilibrium model with New Keynesian features and mortgage debt. He introduces a couple of important improvements: (i) newly issued mortgage debt is subject to *both* loan-to-value (LTV) and payment-to-income (PTI) constraints, and (ii) households have the option to pre-pay their mortgages. These two features amplify the transmission from nominal rates to debt, house prices and output via the constraint switching and front-loading effects, and in addition, even when the shock is transitory, create some different aggregate dynamics depending on the prevailing mortgage contract. In fact, when mortgage rates fall temporarily, borrowers rush to lock in lower rates and take larger loans in the FRM economy, resulting in a larger response of output, debt and house prices. The key to understanding the difference in the response to temporary shocks between these two papers is to recognize that the temporary shock in Greenwald (2018) moves the long end of the yield curve (e.g. via unconventional monetary policies), while the temporary shock in Garriga, Kydland, and Šustek (2021) is a traditional policy rate shock that leaves the long end unchanged and shifts the short end. In fact, our work, which relies heavily on these two papers, shows that in a model similar to Greenwald's (2018) there are no differences between FRM and ARM economies after a temporary monetary policy shock like the one in Garriga, Kydland, and Šustek (2021). In any case, our paper differs from these two papers in several aspects: (i) our focus is on interest rate hikes rather than drops, (ii) in addition to FRM and ARM economies, we also allow for HRM economies given the empirical cross country evidence on the typical interest fixation periods, and (iii) we analyze the interactions between monetary policy transmission and credit limits under different mortgage contract structures.

There are several papers that have empirically documented the average interest fixation period for mortgages across different economies. A prominent example is Badarınza, Campbell, and Ramadorai (2018) who show that there is a vast heterogeneity across countries in the relative popularity of adjustable rate and fixed rate mortgages, and in particular, that the United Kingdom has a low share of FRM relative to other countries, like the United States or Germany. In fact, the mortgage rate fixation period

in the United Kingdom is typically two or five years as shown in Section 3.2. We also show that the most popular mortgage contract in many other countries is neither fixed nor fully adjustable but has an interest fixation period similar to that in the United Kingdom. Several papers, such as Calza, Monacelli, and Stracca (2013), Di Maggio et al. (2017) or Cloyne, Ferreira, and Surico (2020), have also highlighted the importance of the mortgage interest fixation period for the transmission of monetary policy into consumption. Hence, we introduce this type of contractual arrangement into Greenwald's (2018) model and study the monetary policy transmission mechanism when all mortgages are assumed to be hybrid. In addition, refinancing costs for these type of mortgages are extremely expensive and, for example, in the United Kingdom, vary between 5% and 10% of the outstanding loan amount (Best, Cloyne, Ilzetzki, and Kleven 2020). Thus, we switch off the endogenous refinancing channel and assume an exogenous refinancing rate which we calibrate using UK data.

There are also several structural papers that study the aggregate implications of different mortgage contracts beyond conventional FRMs and ARMs. For example, Campbell and Cocco (2003) analyze inflation indexed FRMs from a risk management standpoint; Greenwald, Landvoigt, and Van Nieuwerburgh (2021) study shared appreciation mortgages (SAMs) and show that indexing payments to aggregate house prices generates losses for financial intermediaries that are quantitatively larger than the benefits obtained by borrowers; Guren, Krishnamurthy, and Mcquade (2021) look at a menu of different contracts to analyze which one is better designed to reduce consumption volatility and default; while Eberly and Krishnamurthy (2014) propose a mortgage contract that allows for a one-time costless conversion from FRM to ARM. The latter is the most similar to the mortgage type that we consider, with the distinction that we do not allow for a conversion choice as these fees are prohibitively expensive. Overall, we are not aware of any other paper that has studied how monetary policy transmits to the real economy through the lens of a structural general equilibrium model that features hybrid rate mortgages and different credit limits.

Finally, our paper is also related to the recent literature that analyzes the interactions between monetary policy and mortgage debt limits. For example, Ferrero, Harrison, and Nelson (2023) look at the optimal policy mix between monetary and macro-prudential policies, but only focus on loan-to-value (LTV) limits. Millard, Rubio, and Varadi (2024) also examine the interaction between monetary policy and a combination of macro-prudential limits from a positive standpoint. In this sense, our paper is closely related to theirs as we also consider PTI and LTV limits, but we further focus on the potentially

different impact of these tools based on the mortgage interest fixation period. In fact, given that in our framework not all borrowers are constrained by the same limit, we are able to show how they complement each other in smoothing the responses of economic variables to monetary policy shocks.


Outline. The rest of the paper is structured as follows. In Section 3.2 we present some motivating evidence on mortgage interest fixation periods and show that the vast majority of countries have mortgages with fixation periods that do not exceed ten years. We present a general equilibrium model that takes into account such a mortgage interest rate structure in Section 3.3. Section 3.4 calibrates the model, which is then used as a laboratory to study the effects of monetary policy shocks on consumption in Section 3.5.1, and its interactions with housing tools in Section 3.5.2. Finally, Section 3.6 concludes.

3.2. Mortgage Structure: Is the UK Market that Different?


The structure of mortgage contracts is highly variable across countries and over time. One of its most important aspects is the interest rate schedule applicable over the life of a mortgage loan as these differences influence monetary policy transmission. Most theoretical and empirical studies distinguish between two groups: (i) *fixed rate mortgages (FRM)* in which a nominal interest rate is fixed throughout the life of the mortgage, and (ii) *adjustable rate mortgages (ARM)* in which the mortgage rate varies over the life of the contract according to market conditions. Although some countries mainly rely on one of these two types of contracts, there is a vast heterogeneity in the interest fixation period.

Data from the BIS (2023) shows that indeed some countries such as France or the United States rely exclusively on mortgages with interest fixation periods greater than ten years or on fully variable rates as in the case of Australia or Hong Kong. Nonetheless, these are not the predominant mortgage contracts in many other countries. For example, in Canada, Ireland, the Netherlands, and New Zealand there is a considerable share of mortgages with a fixed-term shorter than five years. These contracts, referred to as *hybrid rate mortgages (HRM)*, start with a fixed rate and switch automatically to the standard variable rate at the end of the fixed-term portion of the contract, unless the borrower chooses to refinance onto a new mortgage contract instead.

The United Kingdom is also characterized by this type of contracts. As of 2023, the majority of new mortgage lending in the United Kingdom was on a fixed-term either



A. Borrowers' exposure to interest rate risk



B. Outstanding loans by rate type and fix period in the United Kingdom

FIGURE 3.2. Interest Fixation Period in Mortgage Loans

NOTE: Panel A shows the share of mortgages broken down by the interest rate fixation period for several countries in the year 2023. Panel B focuses instead on the United Kingdom over the period 2017 to 2024.

for a period of less than two years or for a period of less than five years. The two- and the five-year contracts have also been the most common over time despite some time variation (Figure 3.2, panel B). Regardless of the downward trend from 2017 until 2022, the majority of loans were extended as two-year fixed-term contracts. After 2022, and coinciding with the policy tightening cycle that begun in the United Kingdom in December 2021, there has been a reversal and more five-year fixed term contracts have been issued. In any case, the prevalence of short-term fixed-rate mortgage contracts makes UK households particularly exposed to the risk of unexpected changes in interest rates relative to households in countries with a predominant share of FRM. In what follows, we will embed hybrid rate mortgages into a New-Keynesian model with long term mortgage debt to study the transmission of monetary policy to households.

In addition, the prevalence of shorter-term mortgage contracts also interacts with financial policies and with lenders' own risk assessment practices. For instance, the FCA's Mortgage Conduct of Business (MCOB) includes a requirement that for mortgages where the interest rate may vary within five years, lenders must verify whether the borrower could still afford payments if interest rates were to rise by a minimum of 100 basis points. This requirement is similar to a payment-to-income constraint and hence we explore how these limits interact with monetary policy under different mortgage contracts.

3.3. The Model Economy

This section presents a general equilibrium model of the housing market that builds extensively on Greenwald (2018). That is, the model features two type of households, borrowers and savers, that trade mortgages between each other. Borrowers also make decisions about consumption, labor supply, and the size of newly purchased houses; while savers can also make bond purchases to further insure themselves against aggregate shocks, but have a fixed housing stock. The production side of the economy has the standard New Keynesian features and monetary policy follows a Taylor rule.

We assume in our model that the refinancing rate is exogenous and constant, unlike Greenwald (2018). While endogenous refinancing has proven to be important in the United States – see for example Wong (2019), Beraja et al. (2019) or Eichenbaum, Rebelo, and Wong (2022) – this is not a structural feature of the United Kingdom and many other countries in which HRMs dominate and refinancing before the contract is due implies a substantial cost. Moreover, the refinancing channel in the United States, and probably in other countries as well, is not important during tightening cycles as the refinancing rate is constant when interest rate gaps are negative (Berger et al. 2021).¹

Our primary modeling contribution is to incorporate *hybrid rate mortgages (HRM)* into this framework. We modify the law of motion of promised payments to reflect that HRM mechanically switch from fixed to adjustable rates after T periods and study how the transmission of monetary policy interacts with different housing tools in this environment.

3.3.1. Households

Set-up. There are two types of representative households: *borrowers* and *savers* with measures χ_b and χ_s , respectively. They differ in their preferences. Savers are more patient than borrowers, i.e. $\beta_s > \beta_b$ where β_j is the discount factor of each type of household. They also have different disutility from working η_j to guarantee that they supply the same amount of labor in the steady state. Nonetheless, each agent type maximizes expected lifetime utility over non durable consumption $c_{j,t}$, housing services $h_{j,t}$, and labor supply $n_{j,t}$

¹ The interest rate gap is the difference between the mortgage rate households are paying and the one prevailing in the market at a given point in time.

$$(3.1) \quad \mathbb{E}_t \sum_{k=0}^{\infty} \beta_j^k u(c_{j,t+k}, h_{j,t+k}, n_{j,t+k})$$

where utility takes the separable form

$$(3.2) \quad u(c, h, n) = \log(c) + \xi \log(h) - \eta_j \frac{n^{1+\varphi}}{1+\varphi}.$$

These two types also differ in the composition of their balance sheets. In addition to labor income, which is subject to a proportional tax τ_y that is rebated in the form of lump-sum transfers T_t , households can get resources from different asset classes. In particular, households can trade *one-period nominal bonds*, whose balances are denoted by b_t and have a real return R_t . These bonds are in zero net supply and are used by the monetary authority as a policy instrument. Moreover, we assume that positions in b_t must be non-negative (i.e., they cannot be used for borrowing) and consequently are only traded by savers in equilibrium.

Both agent types also own *housing*, which produces a service flow each period equal to its stock. Homeowners have to pay maintenance costs, which are a constant fraction δ of the value of the house. The total housing stock is fixed $\bar{H}$, as well as the saver demand $h_{s,t} = \tilde{H}_s$, which ensures that the borrower is the marginal buyer. Only prepaying borrowers can adjust their housing holdings as each mortgage loan is linked to a specific house. Hence, the law of motion for the total start-of-period borrower housing is

$$(3.3) \quad h_{b,t} = \rho h_{b,t}^* + (1 - \rho) h_{b,t-1}$$

where ρ is the exogenous fraction of borrowers pre-paying their loans in a given period. This brings us to the most relevant asset in this economy: *mortgages*. Borrowers can trade long-term mortgage debt with savers in equilibrium and their mortgage balance is denoted by m_t . Mortgage debt is issued in the form of fixed-rate perpetuities with coupons that geometrically decay at a rate ν . These loans are also pre-payable and nominal, and consequently, real balances also decay each period at the rate of inflation π_t . As a result, the law of motion of total start of the period debt balances is

$$(3.4) \quad m_t = \rho m_{i,t}^* + (1 - \rho) (1 - \nu) \pi_t^{-1} m_{t-1}$$

where $m_{i,t}^*$ denotes the newly originated loans which are subject to both loan-to-value (LTV) and payment-to-income (PTI) limits at origination

$$(3.5) \quad m_{i,t}^* \leq \theta^{LTV} p_t^h h_{i,t}^*$$

$$(3.6) \quad m_{i,t}^* \leq \frac{\theta^{PTI} w_t n_{i,t} e_{i,t}}{q_t^*}$$

where p_t^h is the price of housing, $h_{i,t}^*$ is the borrower's new house size, w_t is the wage, n_t is labor supply, q_t^* is the coupon of the newly issued mortgage, and $e_{i,t}$ is an idiosyncratic labor productivity shock with c.d.f. Γ_e .² The parameters θ^{LTV} and θ^{PTI} capture the average LTV and PTI limits.

A mortgage contract carries interest that the borrower has to pay to the saver. All future mortgage payments associated with a given loan are subject to a proportional tax $\Delta_{q,t}$, which follows an AR(1) process. Independently of this tax, how mortgage interests are paid depends on the contract type. For adjustable rate mortgages (ARM) these payments vary every period according to the policy rate; while for fixed rate mortgages (FRM), the coupon is fixed at origination and consequently the interest on the stock q_t differs from that on the flow q_t^* . Despite these being the two most studied mortgage contracts, the typical mortgage in many countries, including the United Kingdom, is a hybrid between the two in which interest payments are fixed during T years before they vary according to the policy rate. Denoting by $x_{b,t-1}$ and $x_{s,t-1}$ the total promised payments on existing debt by borrowers and savers, we can specify the evolution of mortgage payments under these *hybrid rate mortgages (HRM)* as follows:

$$(3.7) \quad x_{b,t}^{HRM} = \sum_{\tau=0}^{T-1} \left[((1-\rho)(1-\nu))^\tau \left(\prod_{i=0}^{\tau-1} \pi_{t-i}^{-1} \right) \rho q_{t-\tau}^* m_{t-\tau}^* \right] + \\ + ((1-\rho)(1-\nu))^T \left(\prod_{i=0}^{T-1} \pi_{t-i}^{-1} \right) q_{t-T}^* m_{t-T}$$

$$(3.8) \quad x_{s,t}^{HRM} = \sum_{\tau=0}^{T-1} \left[((1-\rho)(1-\nu))^\tau \left(\prod_{i=0}^{\tau-1} \pi_{t-i}^{-1} \right) \rho (q_{t-\tau}^* - \Delta_{q,t-\tau}) m_{t-\tau}^* \right] + \\ + ((1-\rho)(1-\nu))^T \left(\prod_{i=0}^{T-1} \pi_{t-i}^{-1} \right) (q_{t-T}^* - \Delta_{q,t-T}) m_{t-T}$$

² The labor productivity shock $e_{i,t}$ is used to split borrowers into PTI- and LTV-constrained households. In particular, a fraction $\int^{\bar{e}} e_i d\Gamma_e(e_i)$ is constrained by the PTI limit, while the remaining fraction $1 - \Gamma_e(\bar{e})$ is constrained by the LTV limit.

where the first term corresponds to the fixed part of the contract and the second to the variable rate. That means that borrowers that refinance in period $\tau < T$, and consequently are still on the fixed part of the contract, are paying the rate that was prevailing at the time of refinancing $q_{t-\tau}^*$, while borrowers that has not yet refinance by period T are already in the variable rate.

Borrower's Problem. The borrower chooses consumption $c_{b,t}$, labor supply $n_{b,t}$, the size of newly purchased houses $h_{b,t}^*$, and the face value of newly issued mortgages m_t^* to maximize lifetime utility subject to the borrowing constraints (3.5) - (3.6) and the budget constraint

$$(3.9) \quad c_{b,t} \leq (1 - \tau_y) w_t n_{b,t} + \rho \left(m_t^* - (1 - \nu) \pi_t^{-1} m_{t-1} \right) - \pi_t^{-1} \left((1 - \tau_y) x_{b,t-1} + \nu m_{t-1} \right) - \delta p_t^h h_{b,t-1} - \rho p_t^h \left(h_{b,t}^* - h_{b,t-1} \right) + T_{b,t}$$

where the right hand side is the sum of labor income $(1 - \tau_y) w_t n_{b,t}$, net mortgage issuance $\rho \left(m_t^* - (1 - \nu) \pi_t^{-1} m_{t-1} \right)$ and transfers $T_{b,t}$ minus interests and principal payments $\pi_t^{-1} \left((1 - \tau_y) x_{b,t-1} + \nu m_{t-1} \right)$, net housing purchases $\rho p_t^h \left(h_{b,t}^* - h_{b,t-1} \right)$ and housing maintenance costs $\delta p_t^h h_{b,t-1}$.

Saver's Problem. The saver also choses consumption $c_{s,t}$, labor supply $n_{s,t}$ and the face value of newly issued mortgages m_t^* to maximize lifetime utility subject to the budget constraint

$$(3.10) \quad c_{s,t} \leq (1 - \tau_y) w_t n_{s,t} + \pi_t^{-1} x_{s,t-1} - \rho \left(m_t^* - (1 - \nu) \pi_t^{-1} m_{t-1} \right) - \delta p_t^h \tilde{H}_s - \left(R_t^{-1} b_t - \pi_t^{-1} b_{t-1} \right) + \Pi_t + T_{s,t}$$

where the right hand side is the sum of labor income $(1 - \tau_y) w_t n_{s,t}$, mortgage payments $\pi_t^{-1} x_{s,t-1}$, intermediate profits Π_t , and transfers $T_{s,t}$ minus housing maintenance $\delta p_t^h \tilde{H}_s$, net bond purchases $\left(R_t^{-1} b_t - \pi_t^{-1} b_{t-1} \right)$ and net mortgage issuance $\rho \left(m_t^* - (1 - \nu) \pi_t^{-1} m_{t-1} \right)$.

3.3.2. Production

The production side of the economy has the standard New Keynesian ingredients. A perfectly competitive final good producer that, using intermediate goods $y_t(i)$ as inputs,

produces output Y_t . That is, the final good producer solves the static problem

$$(3.11) \quad \max_{y_t(i)} P_t \left[\int_0^1 y_t(i)^{\frac{\lambda-1}{\lambda}} di \right]^{\frac{\lambda}{\lambda-1}} - \int_0^1 P_t(i) y_t(i) di$$

where P_t is the price of the final good and $P_t(i)$ is the price of each intermediate input $y_t(i)$. The first order condition of this problem gives a demand function for good i

$$(3.12) \quad y_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\lambda} Y_t .$$

To meet this final good producer's demand, the producer of intermediate good i operates a linear production function

$$(3.13) \quad y_t(i) = a_t n_t(i)$$

where $n_t(i)$ is labor hours and a_t is total factor factor productivity (TFP) that evolves according to

$$(3.14) \quad \log a_t = (1 - \phi_a) \mu_a + \phi_a \log a_{t-1} + \varepsilon_{a,t}$$

where $\varepsilon_{a,t}$ is a white noise process. Cost minimization by firm i determines how much labor to hire each period. In particular, a firm i hires workers until the point where wage equals the marginal cost times the marginal product of labor

$$(3.15) \quad W_t = MC_t(i) a_t .$$

Since these producers have some market power, they also set prices. However, they cannot freely adjust them and are subject to price stickiness of the Calvo-Yun form with indexation, which stipulates that each period a fraction $1 - \zeta$ of firms adjust their price to their optimal (flexible) price and the remaining fraction ζ update their price according to the steady state inflation rate.

3.3.3. Monetary authority

Monetary policy is characterized by a Taylor-type rule of the form

$$(3.16) \quad \log R_t = \log \bar{\pi}_t + \phi_r (\log R_{t-1} - \log \bar{\pi}_{t-1}) + \\ + (1 - \phi_r) [(\log R_{ss} - \log \pi_{ss}) + \phi_\pi (\log \pi_t - \log \bar{\pi}_t)] + \log \eta_t$$

where ϕ_r controls the degree of interest rate smoothing, the subscripts “ss” refer to steady state values, $\bar{\pi}_t$ is a time-varying inflation target and η_t is a temporary interest rate shock. $\bar{\pi}_t$ and η_t are defined by

$$(3.17) \quad \log \bar{\pi}_t = (1 - \phi_\pi) \log \pi_{ss} + \phi_\pi \log \bar{\pi}_{t-1} + \varepsilon_{\bar{\pi},t}$$

$$(3.18) \quad \log \eta_t = \phi_\eta \log \eta_{t-1} + \varepsilon_{\eta,t}$$

where $\varepsilon_{\bar{\pi},t}$ and $\varepsilon_{\eta,t}$ are white noise processes which are orthogonal to each other and that we refer to as an *inflation target shock* and an *interest rate shock*, respectively. We include these two types of policy shocks to be able to distinguish between near-permanent and transitory shocks because it has been shown by Garriga, Kydland, and Šustek (2021) the distinction between the two matters when analyzing the responses under different mortgage contracts.

3.3.4. Key equilibrium conditions

The mortgage contract type affects how borrowers pay back savers. We have shown above how one can specify the low of motion of total promised payments on existing mortgage debt such that one is able to capture the typical UK mortgage payment structure. In this section, we present how it affects equilibrium conditions and what it entails for mortgage pricing. We use the ARM and FRM economies as benchmark to simplify the exposition.

The influence of the low of motion of promised payments appears in the borrower’s and saver’s optimality conditions with respect to the face value of newly issued mortgages. The borrower’s optimality of new debt requires

$$(3.19) \quad 1 = \Omega_{b,t}^m + \Omega_{b,t}^x q_t^* + \mu_t$$

where μ_t is the multiplier on borrower’s aggregate credit limit, and $\Omega_{b,t}^m$ and $\Omega_{b,t}^x$ are the marginal continuation *costs* of taking on an additional unit of debt and of promising

an additional dollar of initial payments, respectively. These two marginal continuation values differ based on how interest is paid. In an economy with just FRM contracts, these values are defined as

$$(3.20) \quad \Omega_{b,t}^{m,FRM} = \mathbb{E}_t \left[\Lambda_{t,t+1}^b \pi_{t+1}^{-1} \left(\nu + (1-\nu) \rho + (1-\nu)(1-\rho) \Omega_{b,t+1}^{m,FRM} \right) \right]$$

$$(3.21) \quad \Omega_{b,t}^{x,FRM} = \mathbb{E}_t \left[\Lambda_{t,t+1}^b \pi_{t+1}^{-1} \left((1-\tau_y) + (1-\nu)(1-\rho) \Omega_{b,t+1}^{x,FRM} \right) \right] .$$

while in an economy with ARM contracts, for which promised payments are no longer an endogenous state variable as they change every period, $(\Omega_{b,t}^m + \Omega_{b,t}^x q_t^*)$ can be combined into a single term $\Omega_{b,t}^{ARM}$ that represents the total continuation cost of an additional unit of debt and it is given by

$$(3.22) \quad \Omega_{b,t}^{ARM} = \mathbb{E}_t \left[\Lambda_{t,t+1}^b \pi_{t+1}^{-1} \left((1-\tau_y) q_t^* + \nu + (1-\nu) \rho + (1-\nu)(1-\rho) \Omega_{b,t+1}^{ARM} \right) \right] .$$

Finally, turning to the hybrid rate mortgage (HRM) economy, we see how these marginal continuation values have some similarities with those from the FRM and ARM economies. In fact, the marginal continuation cost of taking an additional unit of debt under HRM is identical to than under FRM, i.e.

$$(3.23) \quad \Omega_{b,t}^m = \Omega_{b,t}^{m,FRM} = \Omega_{b,t}^{m,HRM} .$$

Moreover, the marginal cost of promising an additional dollar of initial payments is identical to that of the FRM up to period T , when the contract mechanically switches to adjustable rates. After period T , it is equal to 0 as for ARM contracts. This is reflected by the finite sum that characterizes this marginal cost in the HRM economy

$$(3.24) \quad \Omega_{b,t}^{x,HRM} = \sum_{\tau=1}^T (1-\rho)^{\tau-1} (1-\nu)^{\tau-1} \mathbb{E}_t \left[\left(\prod_{j=0}^{\tau-1} \Lambda_{t+j,t+j+1}^b \pi_{t+j+1}^{-1} \right) (1-\tau_y) \right]$$

The saver's optimality condition with respect to newly issued mortgage debt is also affected by the low of motion of promised payments. In general, it requires that

$$(3.25) \quad 1 = \Omega_{s,t}^m + \Omega_{s,t}^x (q_t^* - \Delta_{q,t})$$

where $\Omega_{s,t}^m$ and $\Omega_{s,t}^x$ are the marginal continuation *benefits* of an additional unit of debt and of an additional dollar of initial payments, respectively. Similarly to the borrower's marginal costs, these marginal benefits under HRM contracts have some similarities

with their counterparts under FRM and ARM contracts. In fact, it is also the case that the marginal benefit of an additional unit of debt is identical under FRM and HRM contracts. That is,

$$(3.26) \quad \Omega_{s,t}^m = \Omega_{s,t}^{m,FRM} = \Omega_{s,t}^{m,HRM} = \mathbb{E}_t \left[\Lambda_{t,t+1}^s \pi_{t+1}^{-1} \left(\rho (1 - \nu) + (1 - \rho)(1 - \nu) \Omega_{s,t+1}^m \right) \right] .$$

Moreover, the marginal benefit of an additional dollar of initial payments under HRM is also equal to its FRM counterpart up to period T when the contract switches to adjustable rates and it is zero afterwards as in the ARM case as shown below:

$$(3.27) \quad \Omega_{s,t}^{x,HRM} = \sum_{\tau=1}^T (1 - \rho)^{\tau-1} (1 - \nu)^{\tau-1} \mathbb{E}_t \left[\left(\prod_{j=0}^{\tau-1} \Lambda_{t+j+1,t+j}^s \pi_{t+j+1}^{-1} \right) \right] .$$

The marginal continuation benefits of an additional dollar of initial payments in the FRM economy as well as the marginal continuation benefit of an additional unit of debt in the ARM economy are reproduced in Appendix C.1.1.

3.4. Calibration

This section describes the calibration procedure and shows how the model is able to fit the data along several dimensions. The calibration strategy follows the common recipe of setting some parameters externally, while others are chosen jointly with the objective of minimizing the distance between a collection of data and model moments.

We choose the HRM economy with 2 year fixes as the benchmark for calibration. As noted above in Figure 3.2, these are the most common contracts in the United Kingdom and are representative of about 50% of the mortgage loan market. Nonetheless, we also show how steady state moments for a given calibrated parameter vector would change if we were to assume that all mortgage contracts are either on the fixed or the adjustable rate.

3.4.1. Externally calibrated parameters

Demographics & preferences. The fraction of borrowers χ_b is set to match the share of mortgagors whose savings are less than 20% of their total income which is equivalent to 27.74%. This fraction is recovered by combining data from the 2019 English Housing

Survey and Money Dashboard.³ As households have a measure of one, savers represent 72.26% of households in our model economy. Their discount factor β_s is chosen to pin down a ten-year UK gilt yield of 2.5% as they are the only agents that can save in government bonds in our model economy. This results in a value of $\beta_s = 0.998$. Preferences across types also differ in their labor disutility parameters, η_b and η_s , which are chosen to guarantee that borrowers and savers supply the same amount of labor in steady state, $n_{b,ss} = n_{s,ss} = 1/3$. This requires that $\eta_b = 7.518$ and $\eta_s = 5.775$.

The remaining preference parameters that are not internally calibrated are set to standard values in the literature. In particular, the housing utility weight is set to $\xi = 0.25$, while the inverse of the Frisch elasticity is set to $\varphi = 1.0$.

Income process. For the income shock distribution Γ_e , we follow Greenwald (2018) and choose a log-normal specification $\log e_{i,t} \sim \mathcal{N}(-\sigma_e/2, \sigma_e^2)$ which implies that

$$\int_{\bar{e}_t} e_i d\Gamma_e(e_i) = \Phi \left(\frac{\log \bar{e}_t - \sigma_e^2/2}{\sigma_e} \right)$$

where Φ is the CDF of the standard normal distribution. To capture the dispersion in which constraint is binding we set σ_e match the standard deviation of $\log(PTI_{i,t}) - \log(LTV_{i,t})$ in the data. This term is the difference of individual borrowers' log PTI and LTV ratios at origination, which equals $\log e_{i,t}$ in the model. Using the debt service ratio (DSR) as a proxy for the PTI, we find that the UK average of this series is 0.53. Hence, we set σ_e to that value. Labor income is taxed at the rate $\tau_y = 0.212$ which corresponds to the national average prior to interest mortgage deductions in the United Kingdom.

Housing & mortgages. For the debt limit parameters, we set $\theta^{PTI} = 0.36$ and $\theta^{LTV} = 0.85$ which is consistent with the UK empirical distribution of these limits as most mortgagors bunch around those ratios. The amortization parameter is set to the average weighted amortization as a fraction of the loan amount, which in the United Kingdom amounts to 0.57% monthly. Hence, $\nu = 1.71\%$ as one model period corresponds to a quarter. The exogenous refinancing rate is calibrated such that the average duration on a house is ten years, consistently with the UK average. The log housing stock $\log \bar{H}$ is calibrated so

³ The English Housing Survey is used to get homeownership rates for outright owners and mortgagors, while Money Dashboard (MDB) data is useful to compute how many of those have savings below some threshold, consistent with the model assumption about borrowers being constrained by one of the two credit limits.

TABLE 3.1. Parameter values

Parameter	Interpretation	Value	Internal / Jointly
<i>Demographics & preferences</i>			
χ_b	Fraction of borrowers	27.74%	N
β_b	Borr. discount factor	0.957	YY
β_s	Saver discount factor	0.998	Y
ξ	Housing utility weight	0.25	N
η_b	Borr. labor disutility	7.518	Y
η_s	Saver labor disutility	5.775	Y
φ	Inv. Frisch elasticity	1.0	N
<i>Income process</i>			
σ_e	Income dispersion	0.53	N
τ_y	Income tax rate	0.212	N
<i>Housing & mortgages</i>			
θ^{PTI}	Max PTI ratio	0.36	N
θ^{LTV}	Max LTV ratio	0.85	N
ν	Mortgage amortization	1.71%	N
ρ_b	Refinancing rate	0.10	N
δ_h	Housing depreciation	0.005	N
$\log \bar{H}$	Log housing stock	2.256	Y
$\log \bar{H}_s$	Log saver housing stock	1.678	YY
μ_q	Term premium (mean)	0.36%	YY
ϕ_q	Term premium (pers.)	0.852	N
<i>Productive technology</i>			
μ_a	Mean (TFP shock)	1.015	Y
ϕ_a	Persistence (TFP shock)	0.9	N
σ_a	Standard deviation (TFP shock)	0.05	N
λ	Variety elasticity	6.0	N
ζ	Price stickiness	0.75	N
<i>Monetary authority</i>			
ϕ_r	Interest rate smoothing	0.8336	N
φ_π	Taylor rule weight on inflation	1.497	N
π_{ss}	Steady state inflation	1.005	Y
$\phi_{\bar{\pi}}$	Persistence (infl. target shock)	0.994	N
ϕ_η	Persistence (interest rate shock)	0.3	N

NOTE. The model is calibrated at quarterly frequency. Parameters denoted “Y” in the “Internal / Jointly” column are chosen implicitly to match a particular moment at steady state, while those denoted “YY” are chosen jointly to minimize the distance between data and model moments.

that the price of housing is unity at steady state, $p_{ss}^h = 1$. The depreciation rate of the housing stock is set to a standard 0.5% per quarter. Finally, the persistence of the term premium shock is set to the value used by Greenwald (2018).

Productive technology. The persistence and the standard deviation of the TFP shock are taken from COMPASS, the workhorse DSGE model used at the Bank of England for policy analysis and forecasting (Burgess, Fernandez-Corugedo, Groth, Harrison, Monti, Theodoridis, and Waldron 2013). We take their median estimates and set $\rho_a = 0.9$ and $\sigma_a = 0.05$. The mean of the TFP process is chosen such that output in steady state equals one, $y_{ss} = 1$. This results in a value of $\mu_a = 1.015$.

The price-setting parameters take on conventional values: the fraction of firms updating their price ζ is set to 0.75 so that the average price duration is 4 quarters, while the elasticity of substitution λ among varieties is set to 6.0 such that the steady state mark-up $\lambda/(\lambda - 1)$ is 1.2.

Monetary authority. In the Taylor rule, the interest rate smoothing term and the weight on inflation are taken from COMPASS. Using their median estimates we set $\phi_r = 0.8336$ and $\phi_\pi = 1.497$. The mean of the inflation target shock is calibrated to target a 2% inflation rate in the steady state, which results in $\pi_{ss} = 1.005$.

To calibrate the inflation target shock and the interest rate shock processes, we follow Garriga, Kydland, and Šustek (2021) by setting the persistence of the inflation target shock and the transitory interest rate shock to $\varphi_{\bar{\pi}} = 0.994$ and $\varphi_\eta = 0.3$, respectively.⁴

3.4.2. Internally calibrated parameters, targets, and model fit

The remaining three parameters: the borrower discount factor, β_b , the log saver housing stock, $\log \bar{H}_s$, and the mean of the term premium shock, μ_q , are jointly chosen to match three moments in the data: the borrower's house value to income which equals 5.0, the saver's house value to income which is slightly higher and equal to 6.4 and the annualized mortgage rate which is assumed to have a one percentage point spread over the government bond yield.

Table 3.2 shows how the model with HRM and two-year fixes is able to match these moments (top block), as well as to match the steady state targets mentioned in the previous section, after setting $\beta_b = 0.957$, $\log \bar{H}_s = 1.678$ and $\mu_q = 0.36\%$. Moreover, we also

⁴ Using the same parameters for these two shock process helps relating our findings to theirs.

TABLE 3.2. Targets and model fit

Moment	HRM (T=8)	Target / Data	FRM	ARM
<i>Targeted jointly</i>				
House value to income (borr.)	5.08	5.0	4.995	4.945
House value to income (saver)	6.36	6.4	6.388	6.410
Mortgage Rate	3.03%	3.5%	5.33%	5.33%
<i>Steady state targets</i>				
10 Year Gilt Yield	2.5%	2.5%	2.5%	2.5%
Inflation	2.08%	2%	2.08%	2.08%
Output	1.0	$y_{ss} = 1$	1.00	1.00
House price	1.132	$p_{ss}^h = 1$	1.368	1.358
Hours worked (borr.)	0.325	$n_{b,ss} = 1/3$	0.331	0.334
Hours worked (saver)	0.332	$n_{b,ss} = 1/3$	0.330	0.329

NOTE. This table shows the model's ability to capture certain features of the UK economy when calibrated using HRM with 2 years in the fixation period. For comparability, these moments are also shown in the counterfactual economies with FRM or ARM at the same parameter values.

show in the last two columns of Table 3.2 how these steady state moments change when mortgages are assumed to be either fixed or adjustable rate mortgages. Interestingly, the house price and mortgage rate are higher in the FRM and ARM economies.


3.5. Results

3.5.1. Monetary policy pass-through and its effects on consumption


We study the effects of monetary policy shocks on consumption across the three mortgage contract economies. In doing so we distinguish between temporary and persistent shocks as in Garriga, Kydland, and Šustek (2021). These shocks have been calibrated such that they die out after the same number of quarters as in their paper.

Independently of the persistence of the shock, a first step to understand how unexpected movements in the policy rate R_t affect consumption through the housing market is to analyse the monetary policy pass-through to mortgage rates. Figure 3.3 shows how a 1% increase in R_t translates into changes in the average mortgage rate in the economy.⁵

⁵ Since we model long-term mortgage debt, there are two relevant mortgage rates: (i) on newly issued debt, q_t^* , and (ii) on existing debt q_t . The impulse response functions in the right panel of Figure 3.3 shows the average between the two.



A. Temporary Monetary Policy Shock



B. Persistent Inflation Target Shock

FIGURE 3.3. Monetary Policy Pass-Through

NOTE. Responses are normalized such that R_t increases by 1% upon impact in the HRM, FRM & ARM economies. Mortgage rates are expressed as percentage point (annualized) deviations from steady state.

Turning first to the temporary monetary policy shock (panel A), we see that in a *ARM economy* there is a one-to-one pass through and the average mortgage rate response is identical to that of the nominal interest rate. On the other extreme, the average mortgage rate in the *FRM economy* almost does not respond to the increase in the policy rate. As the refinancing rate is rather low, the increase in the mortgage rate on newly issue debt, which is also smaller than in the *ARM economy*, does not get reflected on the overall rate. Finally, in the *HRM economy*, the average mortgage rate increases upon impact but not as much as the increase in R_t , and therefore, the impact on mortgage payments is weaker than in the *ARM economy* but stronger than in the *FRM economy*. This is shown in the subplot in the first row, third column of Figure 3.4.

Despite the differential response of nominal mortgage payments across the three economies, the response of consumption is almost identical. This is explained by the




FIGURE 3.4. Response to a 1% (temporary) monetary policy shock

NOTE. A value of 1 represents a 1% increase relative to the steady state except for mortgage payments, which are measured in percentage points. Output, borrower and saver's consumption are expressed in real terms, and together with inflation and the nominal interest rate are annualized.

larger role of the New Keynesian channel relative to the cash-flow channel when shocks are temporary. Firms respond to the temporarily higher policy rates by decreasing production. This creates downward pressure on prices, but only temporarily as the inflation rate returns to its steady state level within a few quarters. As the inflation rate moves in the opposite direction of the nominal interest rate, the real rate increases even more than the policy rate, and as a result, monetary policy has real effects, as shown by the response of output and consumption that fall substantially in the short run – see bottom panel of Figure 3.4. The drop in consumption is heterogenous across the two types of agents. In fact, borrower's consumption falls almost four times as much as saver's consumption because the latter can partially smooth out the fall in labor income by investing in more profitable bonds. Given that borrowers do not have this margin of adjustment, they try to compensate for increased real mortgage payments by working more hours. However, they cannot fully compensate the fall in wages by their increase in hours, leading to a large drop in consumption. In any case, it is interesting how the response through hours worked is stronger in the *ARM economy* as real mortgage payments increase more – see Figure A1.




FIGURE 3.5. Response to a 1% (persistent) inflation target shock

NOTE. A value of 1 represents a 1% increase relative to the steady state except for mortgage payments, which are measured in percentage points. Output, borrower and saver's consumption are expressed in real terms, and together with inflation and the nominal interest rate are annualized.

The persistent monetary policy shock, modeled as an inflation target shock, has similar implications in terms of the pass-through to mortgage rates. As shown in Panel B of Figure 3.3, the average mortgage rate response in the *ARM economy* is again identical to the evolution of the nominal interest rate. The one-to-one pass through shows up immediately in the response of nominal mortgage payments, which increase persistently and are above their initial steady state value even after ten years – see 1st row, 3rd column of Figure 3.5. On the other hand, the pass-through to mortgage rates in the *FRM economy* takes longer to materialize. Borrowers exogenously refinance into persistently higher rates making the average mortgage rate in the economy increase only gradually. Finally, the pass-through in the *HRM economy* shares some features with the ARM and FRM economies. The average mortgage rate initially increases gradually up to the 8th quarter, when rates switch to being adjustable, and consequently, the average mortgage rate jumps up and then slowly decays given the persistent nature of the shock. As a result, nominal mortgage payments only jump up after two years which gets reflected in borrowers' consumption response.

Unlike for the temporary shock, the persistent change in the policy rate and consequently in mortgage rates has different effects on consumption depending on the mortgage contract structure. These effects are mostly distributional and wash out in the aggregate because firms adjust via prices rather than output, making the rate of inflation increase one-to-one with the nominal interest rate, which in turn leaves the real interest rate unchanged. As result, output and aggregate consumption are not affected and monetary policy has almost no real effects. Nevertheless, there is a redistribution of consumption from borrowers to savers as the persistent nature of the shock makes the increase in mortgage payments more costly to offset. This is particularly important in the *ARM economy* as the increase in mortgage rate is more pronounced and hence the drop in borrower's consumption and the increase in saver's consumption are larger than in the FRM and HRM economies. In other words, the redistributive effects of the shock are amplified when interest fixation periods are shorter.

3.5.1.1. *The length of the fixation period and its impact on consumption*

Until now we have focused on the consumption responses to temporary and persistent monetary policy shocks under the assumption that borrowers spend two years (8 quarters) under a fixed rate before automatically switching to an adjustable rate. This is the typical duration in the United Kingdom, however, the average duration varies across countries as shown in Section 3.2. It is important to recognize that in reality borrowers also have different number of periods left in the fixed part of their mortgage contract when the shock hits. Unfortunately, we cannot speak about the distribution of mortgage interest durations in the cross section as there is no heterogeneity on that front in our model economy. Nonetheless, as a first approximation to the problem we show in Figure 3.6 the response of borrower's and saver's consumption to an inflation target shock in HRM economies with different durations on the fixed part of the mortgage. In particular, the responses are depicted for the ARM, FRM and HRM economies with $T = \{1, \dots, 8\}$. A clear message arises from this figure: the lower T is (lighter lines), the larger is the redistribution from borrower's consumption to saver's consumption since these responses are more and more similar to those in an economy with only ARMs.

The response of borrower's and saver's consumption to the temporary monetary policy shock for different HRM economies is not shown here because the cash-flow channel is unimportant when the shock is transitory and therefore consumption responses are independent of the mortgage contract structure, and consequently, of the duration in the fixed part of the contract for HRMs.




FIGURE 3.6. Consumption response to a 1% (persistent) inflation target shock under different contract durations

NOTE. A value of 1 represents a 1% increase relative to the steady state. The blue and the red lines corresponds to the FRM and the ARM economies. The different shades of black lines represent the response in a HRM economy, with $T=8$ being the darker line and $T=1$ corresponding to the lighter line.

3.5.2. Credit limits and the monetary policy transmission

In this section we focus on the interaction between monetary policy and credit limits. We carry out two types of exercises: (i) we analyze the responses to temporary and persistent shocks under different calibrations for the PTI and LTV limits, and (ii) we explore how the split between LTV- and PTI-constrained borrowers affects the strength of the transmission of monetary policy and its interaction with credit limits. For simplicity, we focus on ARM and FRM economies as the two extremes of HRM contracts and leave out other interest fixation periods from this analysis.

3.5.2.1. Alternative LTV and PTI limits: loose vs. tight credit

To understand the effects of different PTI and LTV calibrations on the transmission of monetary policy into the real economy, we compare the benchmark economy for which the population average credit limits are $\theta_{LTV} = 0.85$ and $\theta_{PTI} = 0.36$, with three counterfactual economies: (i) the *Loose LTV* economy that has a 20% looser maximum LTV limit, i.e. $\theta_{LTV} \approx 1.0$, (ii) the *Loose PTI* economy that has a 20% looser maximum PTI limit, i.e. $\theta_{PTI} \approx 0.43$, and (iii) the *Loose Credit* economy in which both PTI and LTV are 20% looser relative to the benchmark economy.

Figure 3.7 depicts the response to a monetary policy shock that raises the policy rate by 1% upon impact in the *ARM economy* under these four different calibrations. Recall that the interest fixation period was irrelevant in explaining the response to this shock as the New Keynesian channel dominated. Similarly, credit conditions, which




FIGURE 3.7. Loose Credit Limits & the Temporary Monetary Policy Shock – ARM Economy

NOTE. A value of 1 represents a 1% increase relative to the steady state, except for F^{LTV} and new issuance which are expressed in percentage points. New debt issuance is defined as: $\rho(m_t^* - (1 - \nu)\pi_t^{-1}m_{t-1})$, the house price is p_t^h , and the average debt limit $\bar{m}_t$. Output and consumption are reported in real terms.

are obviously related to the housing market, are also not important for explaining the response to this shock and we get almost identical responses regardless of the average credit limits. Intuitively, this result also holds for the *FRM economy* as shown in Panel A of Figure A3.

On the other hand, this is no longer true when the economy is hit by a shock that moves nominal interest rates persistently. Figure 3.8 displays the response to a 1% inflation target shock in the *ARM economy* under these four different calibrations. Intuitively, a looser LTV limit implies a lower steady state fraction of LTV-constrained borrowers (86.33% in the benchmark vs. 77.42% in the *Loose LTV* economy). Since the PTI limit is more sensitive to interest rate changes and there are initially more PTI-constrained borrowers, the average debt limit and the stock of debt fall more in the *Loose LTV* economy. Consequently, the lower housing demand pushes down house prices significantly more in this economy relative to the benchmark. In particular, ten years after the shock hits, house prices fall by 1.61% in the *Loose LTV economy*, while they only fall by 1.06% in the benchmark. On the other hand, a looser PTI calibration has the opposite effect: a larger




FIGURE 3.8. Loose Credit Limits & the Inflation Target Shock – ARM Economy

NOTE. A value of 1 represents a 1% increase relative to the steady state, except for F^{LTV} and new issuance which are expressed in percentage points. New debt issuance is defined as: $\rho(m_t^* - (1 - \nu)\pi_t^{-1}m_{t-1})$, the house price is p_t^h , and the average debt limit $\bar{m}_t$. Output and consumption are reported in real terms.


steady state share of LTV-constrained borrowers (91.24%) implies a lower response of the average debt limit, the stock of debt and house prices despite the increased share of PTI-constrained borrowers associated to higher mortgage rates. Finally, when both credit limits are looser the effects associated with a looser LTV limit dominate as there are fewer steady state LTV-constrained borrowers (82.62%). Hence, the stock of debt, the average debt limit and house prices fall more when both credit limits are loose. In fact, ten years after the shock hits, house prices fall by 1.43% in the *Loose Credit economy*, which is in between the fall in the *Loose LTV* and benchmark economies as the looser PTI limit operates in the opposite direction. These results are consistent with other studies that find that looser credit limits amplify the effects of other shocks when these are permanent (Castellanos, Hannon, and Paz-Pardo 2024). In addition, it is interesting to note that differences in credit conditions also affect the consumption response of borrowers and savers. In particular, these are stronger when LTV limits are looser. In fact, the peak of saver's consumption is 25% larger and the lowest level of borrower consumption is 38% lower when LTVs are loose. On the other hand, looser PTI has only tiny effects on consumption.

In a nutshell, loose LTV limits amplify the effects on house prices and mortgage debt as well as the redistribution of consumption from borrowers to savers associated with persistent movements in the policy rate. Given that PTI limits operate in the opposite direction, coordinating the loosening of LTV limits and PTI limits helps in offsetting the effects of rates on house prices and mortgage debt as the latter act as a backstop. These effects are also present in the *FRM economy*, but they are slightly different quantitatively, as shown in Panel B of Figure A3.


3.5.2.2. *The complementarity between LTV and PTI limits*

In the previous section, we have seen that the split between LTV- and PTI-constrained households is a relevant statistic to understand how credit conditions interact with nominal interest rate hikes. Consequently, we explore in this section how these two limits interact with each other. To do so, we consider two counterfactual economies: (i) the *LTV only economy* which imposes only the LTV constraint, and (ii) the *PTI only economy* which imposes only the PTI constraint. These are just two extremes in which all borrowers are constrained by one of the two limits. In fact, in many other available models, which would serve the same purpose as ours, only one of the two constrains binds (Garriga, Kydland, and Šustek 2021, Millard, Rubio, and Varadi 2024). Nonetheless, we follow Greenwald (2018) in allowing for a split between the two, which not only amplifies the transmission from rates to debt and house prices via the constraint switching effect, but also allow us to show how the initial steady state split between these two affects the transmission of monetary policy and how its strength depends on the interest fixation period.

Figure 3.9 shows the response of new debt issuance, the average debt limit, the fraction of borrowers constrained by the LTV, and the house price to a temporary monetary policy shock in the FRM (panel A) and ARM (panel B) economies with either one or both limits in place. Starting with the *PTI only economy* (dotted lines), we see that the constraint switching effect is switched off as there are no changes in F^{LTV} , however, the initial distribution is such that everyone is PTI-constrained. As PTI limits are more sensitive to interest rate changes and there is a weaker pass-through in the *FRM economy*, the response of new issuance, the average debt limit and the house price is substantially weaker in the *FRM economy*. In fact, in this economy the fall in house prices is 2.36 percentage points lower (9.69% vs. 12.15%), the drop in the average debt limit is 7.84 percentage points lower (7.38% vs. 15.22%) and the fall in debt issuance is half of that in the *ARM economy*. Interestingly, these wide differences are not present in the *LTV only*



A. Fixed Rate Mortgage Economy



B. Adjustable Rate Mortgage Economy

FIGURE 3.9. Constraint Switching & Temporary Monetary Policy Shocks

NOTE. A value of 1 represents a 1% increase relative to the steady state, except for F^{LTV} and new issuance which are expressed in percentage points. New debt issuance is defined as: $\rho(m_t^* - (1 - \nu)\pi_t^{-1}m_{t-1})$, the average debt limit: $\bar{m}_t$, and the house price: p_t^h .

economy (dashed lines) in which the steady state distribution is such that everyone is LTV-constrained. Intuitively, LTVs are not affected directly by changes in mortgage rates and hence the interest fixation period and its impact on the pass-through to mortgage rates becomes unimportant to explain the response to the shock.

Finally, turning to the benchmark economy with both LTV and PTI limits in place (solid lines), we see that the fraction of households constrained by the LTV decreases in both ARM and FRM economies as the increase in rates tightens the PTI limit. The stronger pass-through in the *ARM economy* and the high sensitivity of the PTI to changes in mortgage rates results into a larger drop in the fraction of LTV-constrained borrowers F^{LTV} , which is almost four times as large than the fall in the FRM economy. Despite this large difference, we only see a 1.17 percentage points larger drop upon impact on

house prices (4.67% vs. 5.84%), a 1.88 percentage points larger fall in the average debt limit at time 0 (5.0% vs. 6.88%) and nearly no differences in the initial drop of new debt issuance between the *ARM economy* and *FRM economy*. In a nutshell, this implies that the constraint switching effect is more important when the fixation period is shorter, but it is the steady state distribution of LTV- and PTI-constrained households that matters the most.

3.6. Conclusion

In this paper, we provide a structural analysis of the impact of mortgage interest fixation periods on the strength of monetary policy and its interaction with borrower-based macro-prudential limits. We base our analysis in a standard housing model with long term debt and New Keynesian features to which we add hybrid rate mortgages to reflect the various interest fixation periods observed in cross country data.

The distinction between temporary and persistent increases in nominal interest rates as well as the inclusion of two types of credit limits, LTVs and PTIs, are crucial to understand our results. We show that a temporary increase in nominal interest rates leads to increased saving and reduced investment as firms react through changes in production rather than prices. As a result, households reduce consumption temporarily. This New Keynesian channel dominates the housing market channel, and consequently, the interest fixation period and credit conditions on borrowers' mortgages play a minor role in shaping the responses to this shock. On the other hand, shocks that lead to persistent increases in nominal interest rates have different impacts depending on both credit conditions and the interest fixation periods. In particular, an inflation target shock that persistently moves nominal interest rates leads to a redistribution from borrowers to savers which is stronger when LTV limits are looser and interest fixation periods are shorter. Moreover, we also find that the split between LTV and PTI constrained households is essential to get to these results. Otherwise, the highly sensitive PTI limits will amplify the effects of monetary policy shocks, even when temporary, and especially for economies dominated by short interest fixation periods.

Our paper highlights the importance of the interest fixation periods for the conduct of monetary and macro-prudential policies. These results open interesting avenues for future research. For instance, we need further understanding on the time variation of interest fixation periods in the data and how they may be affected by macroeconomic conditions. Another interesting line of research would be to disentangle why there are

such institutional differences across countries in terms of fixation periods and whether these are related to borrowers' preferences or other factors. This may help explain why monetary policy is more powerful in some economies than others.

Bibliography

Aastveit, Knut Are, and André K Anundsen. 2022. "Asymmetric effects of monetary policy in regional housing markets." *American Economic Journal: Macroeconomics* 14 (4): 499–529.

Acharya, Viral V., Katharina Bergant, Matteo Crosignani, Tim Eisert, and Fergal McCann. 2022. "The Anatomy of the Transmission of Macroprudential Policies." *The Journal of Finance* 77 (5): 2533–2575.

Aladangady, Aditya. 2017. "Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata." *American Economic Review* 107 (11): 3415–3446.

Amaral, Francisco, Martin Dohmen, Sebastian Kohl, and Moritz Schularick. 2024. "Interest rates and the spatial polarization of housing markets." *American Economic Review: Insights* 6 (1): 89–104.

Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme. 2017. "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework." *Econometrica* 85 (3): 693–734.

Auclert, Adrien. 2019. "Monetary Policy and the Redistribution Channel." *American Economic Review* 109 (6): 2333–2367.

Badarinza, Cristian, John Y Campbell, and Tarun Ramadorai. 2018. "What calls to ARMs? International evidence on interest rates and the choice of adjustable-rate mortgages." *Management Science* 64 (5): 2275–2288.

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra. 2019. "Regional Heterogeneity and the Refinancing Channel of Monetary Policy." *The Quarterly Journal of Economics* 134 (1): 109–183.

- Berger, David, Konstantin Milbradt, Fabrice Tourre, and Joseph Vavra. 2021. "Mortgage Prepayment and Path-Dependent Effects of Monetary Policy." *American Economic Review* 111 (9): 2829–2878.
- Bernanke, Ben S., and Alan S. Blinder. 1992. "The Federal Funds Rate and the Channels of Monetary Transmission." *The American Economic Review* 82 (4): 901–921. Publisher: American Economic Association.
- Bernanke, Ben S, and Mark Gertler. 1995. "Inside the Black Box: The Credit Channel of Monetary Policy Transmission." *Journal of Economic Perspectives* 9 (4): 27–48.
- Best, Michael Carlos, James S Cloyne, Ethan Ilzetzki, and Henrik J Kleven. 2020. "Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches." *The Review of Economic Studies* 87 (2): 656–690.
- Bianchi, Francesco, Cosmin Ilut, and Hikaru Saijo. 2024. "Diagnostic Business Cycles." *Review of Economic Studies* 91 (1): 129–162.
- BIS. 2023. "Macroprudential policies to mitigate housing market risks." CGFS Papers No 69.
- Blanchard, Olivier, and Roberto Perotti. 2002. "An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output." *The Quarterly Journal of Economics* 117 (4): 1329–1368.
- Burgess, Stephen, Emilio Fernandez-Corugedo, Charlotta Groth, Richard Harrison, Francesca Monti, Konstantinos Theodoridis, and Matt Waldron. 2013. "The Bank of England's Forecasting Platform: COMPASS, MAPS, EASE and the Suite of Models." SSRN Electronic Journal.
- Calvo, Guillermo A. 1983. "Staggered prices in a utility-maximizing framework." *Journal of Monetary Economics* 12 (3): 383–398.
- Calza, Alessandro, Tommaso Monacelli, and Livio Stracca. 2013. "Housing Finance and Monetary Policy." *Journal of the European Economic Association* 11 (s1): 101–122.
- Campbell, John Y., and Joao F. Cocco. 2003. "Household Risk Management and Optimal Mortgage Choice." *Quarterly Journal of Economics* 118 (4): 1449–1494.

- Canova, Fabio, and Luca Sala. 2009. "Back to square one: Identification issues in DSGE models." *Journal of Monetary Economics* 56 (4): 431–449.
- Castellanos, Juan, Andrew Hannon, and Gonzalo Paz-Pardo. 2024. "The aggregate and distributional implications of credit shocks on housing and rental markets." SSRN Electronic Journal.
- Chari, V.V., Patrick J. Kehoe, and Ellen R. McGrattan. 2008. "Are structural VARs with long-run restrictions useful in developing business cycle theory?" *Journal of Monetary Economics* 55 (8): 1337–1352.
- Christiano, Lawrence, Martin Eichenbaum, and Charles Evans. 1999. "Monetary Policy Shocks: What Have We Learned and to What End?" In *Handbook of Macroeconomics*, edited by John B. Taylor and M. Woodford, vol. 1A, chap. 2, 65–148: Elsevier.
- Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt. 2016. "Unemployment and Business Cycles." *Econometrica* 84 (4): 1523–1569.
- Christiano, Lawrence J., Martin Eichenbaum, Robert Vigfusson, Patrick J. Kehoe, and Mark W. Watson. 2006. "Assessing Structural VARs [with Comments and Discussion]." *NBER Macroeconomics Annual* 21: 1–105.
- Christiano, Lawrence J., Mathias Trabandt, and Karl Walentin. 2010. "DSGE Models for Monetary Policy Analysis." In *Handbook of Monetary Economics*, vol. 3, 285–367: Elsevier.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113 (1): 1–45.
- Cloyne, James, Clodomiro Ferreira, and Paolo Surico. 2020. "Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism." *The Review of Economic Studies* 87 (1): 102–129.
- Cloyne, James, Kilian Huber, Ethan Ilzetzki, and Henrik Kleven. 2019. "The Effect of House Prices on Household Borrowing: A New Approach." *American Economic Review* 109 (6): 2104–2136.

- Creel, Michael, and Dennis Kristensen. 2011. "Indirect Likelihood Inference." Barcelona GSE Working Paper Series, No. 558.
- De Nardi, Mariacristina, Giulio Fella, and Gonzalo Paz-Pardo. 2020. "Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare." *Journal of the European Economic Association* 18 (2): 890–926.
- De Nardi, Mariacristina, Giulio Fella, and Gonzalo Paz-Pardo. 2024. "Wage risk and government and spousal insurance." *Review of Economic Studies*: forthcoming.
- Di Maggio, Marco, Amir Kermani, Benjamin J. Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao. 2017. "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging." *American Economic Review* 107 (11): 3550–3588.
- Dias, Daniel A, and João B Duarte. 2019. "Monetary policy, housing rents, and inflation dynamics." *Journal of Applied Econometrics* 34 (5): 673–687.
- Dias, Daniel A, and Joao B Duarte. 2022. "Monetary Policy and Homeownership: Empirical Evidence, Theory, and Policy Implications.", Working Paper, Board of Governors of the Federal Reserve System (US).
- Eberly, Janice, and Arvind Krishnamurthy. 2014. "Efficient Credit Policies in a Housing Debt Crisis." *Brookings Papers on Economic Activity*: 73–118.
- Eichenbaum, Martin, Sergio Rebelo, and Arlene Wong. 2022. "State-Dependent Effects of Monetary Policy: The Refinancing Channel." *American Economic Review* 112 (3): 721–761.
- Farhi, Emmanuel, and Iván Werning. 2016. "A Theory of Macprudential Policies in the Presence of Nominal Rigidities." *Econometrica* 84 (5): 1645–1704.
- Favilukis, Jack, Sydney C Ludvigson, and Stijn Van Nieuwerburgh. 2017. "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium." *Journal of Political Economy* 125 (1): 140–223.
- Favilukis, Jack, Pierre Mabilie, and Stijn Van Nieuwerburgh. 2022. "Affordable Housing and City Welfare." *The Review of Economic Studies* 90 (1): 293–330.

- Fella, Giulio. 2014. "A generalized endogenous grid method for non-smooth and non-concave problems." *Review of Economic Dynamics* 17 (2): 329–344.
- Fernández-Villaverde, Jesús, Juan F Rubio-Ramírez, Thomas J Sargent, and Mark W Watson. 2007. "ABCs (and Ds) of Understanding VARs." *American Economic Review* 97 (3): 1021–1026.
- Ferrero, Andrea, Richard Harrison, and Benjamin Nelson. 2023. "House Price Dynamics, Optimal LTV Limits and the Liquidity Trap." *The Review of Economic Studies*: 1–32.
- Floetotto, Max, Michael Kirker, and Johannes Stroebel. 2016. "Government intervention in the housing market: Who wins, who loses?" *Journal of Monetary Economics* 80: 106–123.
- Garriga, Carlos, and Aaron Hedlund. 2020. "Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession." *American Economic Review* 110 (6): 1603–1634.
- Garriga, Carlos, Finn E. Kydland, and Roman Šustek. 2017. "Mortgages and Monetary Policy." *The Review of Financial Studies* 30 (10): 3337–3375.
- Garriga, Carlos, Finn E. Kydland, and Roman Šustek. 2021. "MoNK: Mortgages in a New-Keynesian model." *Journal of Economic Dynamics and Control* 123: 104059.
- Gete, Pedro, and Michael Reher. 2018. "Mortgage Supply and Housing Rents." *The Review of Financial Studies* 31 (12): 4884–4911.
- Gourieroux, C., A. Monfort, and E. Renault. 1993. "Indirect Inference." *Journal of Applied Econometrics* 8: S85–S118.
- Greenwald, Daniel L. 2018. "The Mortgage Credit Channel of Macroeconomic Transmission." SSRN Electronic Journal.
- Greenwald, Daniel L, and Adam Guren. 2024. "Do Credit Conditions Move House Prices?". SSRN Electronic Journal.
- Greenwald, Daniel L., Tim Landvoigt, and Stijn Van Nieuwerburgh. 2021. "Financial Fragility with SAM?" *The Journal of Finance* 76 (2): 651–706.

- Guren, Adam M., Arvind Krishnamurthy, and Timothy J. McQuade. 2021. "Mortgage Design in an Equilibrium Model of the Housing Market." *The Journal of Finance* 76 (1): 113–168.
- Hall, Alastair R., Atsushi Inoue, James M. Nason, and Barbara Rossi. 2012. "Information criteria for impulse response function matching estimation of DSGE models." *Journal of Econometrics* 170 (2): 499–518.
- Hannon, Andrew. 2023. "Falling Behind: Delinquency and Foreclosure in a Housing Crisis." Manuscript.
- Herbst, Edward P., and Benjamin K. Johannsen. 2023. "Bias in Local Projections." Manuscript.
- Iacoviello, Matteo. 2005. "House prices, borrowing constraints, and monetary policy in the business cycle." *American Economic Review* 95 (3): 739–764.
- Iacoviello, Matteo. 2011. "Housing Wealth and Consumption." Board of Governors of the Federal Reserve System - International Finance Discussion Papers.
- Inoue, Atsushi, Chun-Hung Kuo, and Barbara Rossi. 2020. "Identifying the sources of model misspecification." *Journal of Monetary Economics* 110: 1–18.
- Ireland's Department of Finance. 2019. "Institutional Investment in the Housing Market."
- Iskhakov, Fedor, Thomas H. Jørgensen, John Rust, and Bertel Schjerning. 2017. "The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks." *Quantitative Economics* 8 (2): 317–365.
- Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *The American Economic Review* 95 (1): 161–182.
- Jordà, Òscar, and Sharon Kozicki. 2011. "Estimation and Inference by the Method of Projection Minimum Distance: An Application to the New Keynesian Hybrid Phillips Curve*." *International Economic Review* 52 (2): 461–487.
- Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti. 2019. "Credit Supply and the Housing Boom." *Journal of Political Economy* 127 (3): 1317–1350.

- Kaplan, Greg, Kurt Mitman, and Giovanni L Violante. 2020. “The Housing Boom and Bust: Model Meets Evidence.” *Journal of Political Economy* 128 (9): 3285–3345.
- Kelly, Robert, Fergal McCann, and Conor O’Toole. 2018. “Credit conditions, macroprudential policy and house prices.” *Journal of Housing Economics* 41: 153–167.
- Kilian, Lutz. 1998. “Small-sample Confidence Intervals for Impulse Response Functions.” *Review of Economics and Statistics* 80 (2): 218–230.
- Lambertini, Luisa, Caterina Mendicino, and Maria Teresa Punzi. 2013. “Leaning against boom–bust cycles in credit and housing prices.” *Journal of Economic Dynamics and Control* 37 (8): 1500–1522.
- Levy, Antoine. 2022. “Housing Policy with Home-Biased Landlords: Evidence from French Rental Markets.” Manuscript.
- Li, Dake, Mikkel Plagborg-Møller, and Christian K. Wolf. 2023. “Local Projections vs. VARs: Lessons From Thousands of DGPs.” NBER Working Paper No. W30207.
- Lyons, Ronan C. 2018. “Credit conditions and the housing price ratio: Evidence from Ireland’s boom and bust.” *Journal of Housing Economics* 42: 84–96.
- Millard, Stephen, Margarita Rubio, and Alexandra Varadi. 2024. “The Macroprudential Toolkit: Effectiveness and Interactions.” *Oxford Bulletin of Economics and Statistics* 86 (2): 335–384.
- Muñoz, Manuel A, and Frank Smets. 2022. “Macroprudential policy and the role of institutional investors in housing markets.” ECB Working Paper Series.
- Nakajima, Makoto, and Irina A Telyukova. 2020. “Home equity in retirement.” *International Economic Review* 61 (2): 573–616.
- Olea, José Luis Montiel, Mikkel Plagborg-Møller, Eric Qian, and Christian K Wolf. 2024. “Double Robustness of Local Projections and Some Unpleasant VARithmetic.” NBER Working Paper No. 32495.
- Oliveira, João G, and Leonor Queiró. 2023. “Mortgage borrowing caps: leverage, default and welfare.” Bank of Portugal Working Papers.

- Oosthuizen, Dick. 2023. "Institutional Housing Investors and the Great Recession." Manuscript.
- Paz-Pardo, Gonzalo. 2024. "Homeownership and Portfolio Choice Over the Generations." *American Economic Journal: Macroeconomics* 16 (1): 207–237.
- Pew Research. 2021. "Who rents and who owns in the U.S."
- Piazzesi, Monika, and Martin Schneider. 2016. "Housing and macroeconomics." *Handbook of Macroeconomics* 2: 1547–1640.
- Plagborg-Møller, Mikkel, and Christian K. Wolf. 2021. "Local Projections and VARs Estimate the Same Impulse Responses." *Econometrica* 89 (2): 955–980.
- Pope, Alun Lloyd. 1990. "Biases of estimators in multivariate non-Gaussian autoregressions." *Journal of Time Series Analysis* 11 (3): 249–258.
- Poterba, James M. 1984. "Tax subsidies to owner-occupied housing: an asset-market approach." *The Quarterly Journal of Economics* 99 (4): 729–752.
- Ramey, V. A. 2016. "Macroeconomic Shocks and Their Propagation." In *Handbook of Macroeconomics*, edited by John B. Taylor and Harald Uhlig, chap. 2, 71–162: Elsevier.
- Ramey, Valerie A. 2011. "Identifying Government Spending Shocks: It's all in the timing." *The Quarterly Journal of Economics* 126 (1): 1–50.
- Romer, Christina D., and David H. Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *The American Economic Review* 94 (4): 1055–1084.
- Rotemberg, Shahar, and Joseph B Steinberg. 2024. "Mortgage interest deductions? Not a bad idea after all." *Journal of Monetary Economics* 144.
- Rotemberg, Julio, and Michael Woodford. 1997. "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy." In *NBER Macroeconomics Annual*, edited by Julio Rotemberg and Ben S. Bernanke, chap. 12, 297–361: MIT Press.
- Ruge-Murcia, Francisco. 2012. "Estimating nonlinear DSGE models by the simulated method of moments: With an application to business cycles." *Journal of Economic Dynamics and Control* 36 (6): 914–938.

- Ruge-Murcia, Francisco. 2020. "Estimating nonlinear dynamic equilibrium models by matching impulse responses." *Economics Letters* 197 (109624).
- Ruge-Murcia, Francisco J. 2007. "Methods to estimate dynamic stochastic general equilibrium models." *Journal of Economic Dynamics and Control* 31 (8): 2599–2636.
- Scalone, Valerio. 2018. "Estimating Non-Linear DSGEs with the Approximate Bayesian Computation: an application to the Zero Lower Bound." Banque de France Working Paper, Number: 688.
- Smets, Frank, and Rafael Wouters. 2003. "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association* 1 (5): 1123–1175.
- Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (3): 586–606.
- Smith, A. A. 1993. "Estimating Nonlinear Time-Series Models Using Simulated Vector Autoregressions." *Journal of Applied Econometrics* 8: S63–S84.
- Sommer, Kamila, and Paul Sullivan. 2018. "Implications of US tax policy for house prices, rents, and homeownership." *American Economic Review* 108 (2): 241–74.
- Stock, James H., and Mark W. Watson. 2018. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments." *The Economic Journal* 128 (610): 917–948.
- Van Bakkum, Sjoerd, Rustom M Irani, Marc Gabarro, and José Luis Peydró. 2023. "Take It to the Limit? The Effects of Household Leverage Caps." SSRN Electronic Journal.
- Wong, Arlene. 2019. "Refinancing and The Transmission of Monetary Policy to Consumption." Working paper, Princeton University.

Appendix A

Appendix to Chapter 1

Appendix A.1. The Smets-Wouters Model

The log-linearized equilibrium conditions of the Smets and Wouters (2007) model take the following form:

$$(A1) \quad \hat{y}_t = c_y \hat{c}_t + i_y \hat{i}_t + z_y \hat{z}_t + \varepsilon_t^g$$

$$(A2) \quad \begin{aligned} \hat{c}_t = & \frac{h/\gamma}{1+h/\gamma} \hat{c}_{t-1} + \frac{1}{1+h/\gamma} \mathbb{E}_t \hat{c}_{t+1} + \frac{wl_c(\sigma_c-1)}{\sigma_c(1+h/\gamma)} \left(\hat{l}_t - \mathbb{E}_t \hat{l}_{t+1} \right) + \\ & - \frac{1-h/\gamma}{(1+h/\gamma)\sigma_c} (\hat{r}_t - \mathbb{E}_t \hat{\pi}_{t+1}) - \frac{1-h/\gamma}{(1+h/\gamma)\sigma_c} \varepsilon_t^b \end{aligned}$$

$$(A3) \quad \hat{i}_t = \frac{1}{1+\beta\gamma^{(1-\sigma_c)}} \hat{i}_{t-1} + \frac{\beta\gamma^{(1-\sigma_c)}}{1+\beta\gamma^{(1-\sigma_c)}} \mathbb{E}_t \hat{i}_{t+1} + \frac{1}{\varphi\gamma^2(1+\beta\gamma^{(1-\sigma_c)})} \hat{q}_t + \varepsilon_t^i$$

$$(A4) \quad \hat{q}_t = \beta(1-\delta)\gamma^{-\sigma_c} \mathbb{E}_t \hat{q}_{t+1} - \hat{r}_t + \mathbb{E}_t \hat{\pi}_{t+1} + (1-\beta(1-\delta)\gamma^{-\sigma_c}) \mathbb{E}_t \hat{r}_{t+1}^k - \varepsilon_t^b$$

$$(A5) \quad \hat{y}_t = \Phi \left(\alpha \hat{k}_t^s + (1-\alpha) \hat{l}_t + \varepsilon_t^a \right)$$

$$(A6) \quad \hat{k}_t^s = \hat{k}_{t-1} + \hat{z}_t$$

$$(A7) \quad \hat{z}_t = \frac{1-\psi}{\psi} \hat{r}_t^k$$

$$(A8) \quad \hat{k}_t = \frac{(1-\delta)}{\gamma} \hat{k}_{t-1} + (1-(1-\delta)/\gamma) \hat{i}_t + (1-(1-\delta)/\gamma) \varphi\gamma^2 \left(1 + \beta\gamma^{(1-\sigma_c)} \right) \varepsilon_t^i$$

$$(A9) \quad \hat{\mu}_t^p = \alpha \left(\hat{k}_t^s - \hat{l}_t \right) - \hat{w}_t + \varepsilon_t^a$$

$$(A10) \quad \hat{\pi}_t = \frac{\beta\gamma^{(1-\sigma_c)}}{1 + \iota_p\beta\gamma^{(1-\sigma_c)}} \mathbb{E}_t \hat{\pi}_{t+1} + \frac{\iota_p}{1 + \beta\gamma^{(1-\sigma_c)}} \hat{\pi}_{t-1} +$$

$$- \frac{\left(1 - \beta\gamma^{(1-\sigma_c)}\xi_p\right) (1 - \xi_p)}{(1 + \iota_p\beta\gamma^{(1-\sigma_c)}) (1 + (\Phi - 1)\varepsilon_p) \xi_p} \hat{\mu}_t^p + \varepsilon_t^p$$

$$(A11) \quad \hat{r}_t^k = \hat{l}_t + \hat{w}_t - \hat{k}_t^s$$

$$(A12) \quad \hat{\mu}_t^w = \hat{w}_t - \sigma_l \hat{l}_t - \frac{1}{1 - h/\gamma} (\hat{c}_t - h/\gamma \hat{c}_{t-1})$$

$$(A13) \quad \hat{w}_t = \frac{\beta\gamma^{(1-\sigma_c)}}{1 + \beta\gamma^{(1-\sigma_c)}} (\mathbb{E}_t \hat{w}_{t+1} + \mathbb{E}_t \hat{\pi}_{t+1}) + \frac{1}{1 + \beta\gamma^{(1-\sigma_c)}} (\hat{w}_{t-1} - \iota_w \hat{\pi}_{t-1}) +$$

$$- \frac{1 + \beta\gamma^{(1-\sigma_c)}\iota_w}{1 + \beta\gamma^{(1-\sigma_c)}} \hat{\pi}_t - \frac{\left(1 - \beta\gamma^{(1-\sigma_c)}\xi_w\right) (1 - \xi_w)}{(1 + \beta\gamma^{(1-\sigma_c)}) (1 + (\lambda_w - 1)\varepsilon_w) \xi_w} \hat{\mu}_t^w + \varepsilon_t^u$$

$$(A14) \quad \hat{r}_t = \rho \hat{r}_{t-1} + (1 - \rho) (r_\pi \hat{\pi}_t + r_y (\hat{y}_t - \hat{y}_t^*)) + r_{\Delta y} ((\hat{y}_t - \hat{y}_t^*) - (\hat{y}_{t-1} - \hat{y}_{t-1}^*)) + \varepsilon_t^r$$

And the (seven) exogenous shocks evolve according to:

$$(A15) \quad \varepsilon_t^a = \rho_a \varepsilon_{t-1}^a + \eta_t^a$$

$$(A16) \quad \varepsilon_t^b = \rho_b \varepsilon_{t-1}^b + \eta_t^b$$

$$(A17) \quad \varepsilon_t^g = \rho_g \varepsilon_{t-1}^g + \rho_{ga} \eta_t^a + \eta_t^g$$

$$(A18) \quad \varepsilon_t^i = \rho_i \varepsilon_{t-1}^i + \eta_t^i$$

$$(A19) \quad \varepsilon_t^r = \rho_r \varepsilon_{t-1}^r + \eta_t^r$$

$$(A20) \quad \varepsilon_t^p = \rho_p \varepsilon_{t-1}^p + \eta_t^p$$

$$(A21) \quad \varepsilon_t^w = \rho_w \varepsilon_{t-1}^w + \eta_t^w$$

Appendix A.2. Data Moments / Targets

A selection of the impulse responses used as targets or data moments in estimation are depicted below. These are obtained after simulating the model at the true parameter vector Θ^* for each of the $S = 100$ draws of the shocks under different identification strategies and using either LPs or SVAR methods for estimation.

A.2.1. Observed Innovation

A.2.1.1. The bias-variance trade off

Figure A1 depicts the response of output, consumption, investment and hours worked to one standard deviation of the monetary policy shock. The dashed line in both panels is the structural IRF that one aims to estimate using either the LP (panel A) or the SVAR (panel B) models. These are depicted with a fan chart to capture the distribution over the different draws of the shock. The median response is plotted with a solid line. In both cases, the sample size is $T = 300$ and the lag length is set to $p = 4$.

From this simple plotting exercise, one learns that the median LP estimated response (solid line in panel A) is very similar to the structural IRF, while the the median estimated




FIGURE A1. Responses to an observed monetary innovation

NOTE. This figure shows the distribution of the estimated responses of output, consumption, investment and hours worked to a monetary innovation that have been estimated using either a LP (panel A) or SVAR (panel B) approach and $p = 4$ lags. The solid line is the median response, while the dash line is the structural IRF.

SVAR response (solid line in panel B) differs substantially at long horizons. Moreover, the distribution of the LP responses is wider than that of SVAR responses as the latter tend to die out at long horizons. In other words, LP has lower bias than SVAR responses, but it comes at the cost of having also a larger variance than SVARs. This result is consistent with the findings in Li, Plagborg-Møller, and Wolf (2023) and it also present in the response to other shocks within the Smets and Wouters model, such as technology or fiscal policy shocks.

A.2.1.2. *Observed innovation vs. observed shock*

It is common knowledge that using the innovation or the shock itself gives the same impulse responses as long as the shocks are independent and identically distributed. Hence, in the context of the Smets and Wouters (2007) model estimating the responses to technology and monetary policy using the shock, i.e. by setting $\tilde{x}_t$ to ε_t^m or ε_t^a , will give the same answer as to using the innovation itself, i.e. setting $\tilde{x}_t$ to η_t^m or η_t^a . Therefore, the results in Section 1.5.1 can be also interpreted as if the econometrician were to observe the shock, but with one caveat. Government spending is correlated with the technology shock, as shown in equation (A17), and hence the response to the innovation is not identical to the response of the shock. I will explore the difference between the innovation and the shock when I study the case in which the econometrician observes a noisy measure of the shock of interest – see Section 1.3.4.3 for a discussion and Section 1.5.3 for the results.

A.2.2. Recursive Identification

As discussed in Section 1.3.4.2, there are two widely used alternatives to identify the shocks through imposing recursive zero restrictions on contemporaneous coefficients. The first one assumes that the policy variable does not respond within the period to other exogenous variables, while the second one imposes that other endogenous variables do not respond to the policy shock within the period. See Ramey (2016) for details.

A.2.2.1. *Technology shock*

The technology shock governs the evolution of TFP in the Smets and Wouters (2007) model. The TFP process follows an AR(1) in logs and it is completely exogenous, as shown in equation (A15). Hence, it is reasonable to assume that the policy variable, TFP, does not respond to other exogenous variables within the period. In fact, that is the




FIGURE A2. Responses to a recursive technology shock

NOTE. This figure shows the distribution of the estimated responses of output, consumption, investment and hours worked to a technology shock identified recursively and that has been estimated using either a LP (panel A) or SVAR (panel B) approach with $p = 4$ lags. The solid line is the median response, while the dash line is the structural IRF and the dotted line is the population LP/SVAR response with infinite lags.

correct assumption as illustrated by the fact that the population response (dotted line) and the structural IRF (dash line) coincide at all horizons. Therefore, the estimated SVAR-IRFs, which rely on a VAR where I order TFP as the first variable, as well as the LP-IRFs, that set $\tilde{x}_t = \varepsilon_t^a$, coincide with the estimated IRFs under the observed innovation assumption. Figure A2 depicts the distribution of the output, consumption, investment and hours worked estimated responses using the aforementioned recursive identification strategy with $T = 300$ and $p = 4$, and in fact, they are identical to the distribution of responses estimated under the observed shock assumption.

Finally, note that the bias-variance trade off is also present here as well as the small sample bias. These issues concern the estimation approach and are independent of the identification strategy.

A.2.2.2. Monetary policy shock

For the monetary policy shock I assume instead that other endogenous variables do not respond to the policy shock within the period as it commonly assumed in the literature, see for example Bernanke and Blinder (1992) or Christiano, Eichenbaum, and Evans (2005). Differently from the technology shock, this assumption does not hold within the




FIGURE A3. Responses to a recursive monetary policy shock

NOTE. This figure shows the distribution of the estimated responses of output, consumption, investment and hours worked to a monetary policy shock identified recursively and that has been estimated using either a LP (panel A) or SVAR (panel B) approach with $p = 4$ lags. The solid line is the median response, while the dash line is the structural IRF and the dotted line is the population LP/SVAR response with infinite lags.

Smets and Wouters (2007) model. This can be seen graphically in Figure A3 in which the population response (dotted line) disagrees with structural IRF (dash line). In fact, real variables respond contemporaneously to a monetary policy shock in the Smets and Wouters (2007) model, which is ruled out by our identification assumption. The distribution of these estimated responses by either LP or SVAR is also plotted in this figure and features the usual bias variance trade off with respect to the population responses.

A.2.3. Direct measures of the shocks of interest

A.2.3.1. Uncorrelated external proxies

Figure A4 plots the distribution of estimated responses to a technology shock under the assumption that the econometrician observes a proxy for the shock and that the noise in the proxy is uncorrelated with other shocks (classical measurement error). As shown from the difference between the dashed and the dotted lines, the presence of uncorrelated measurement error lead to attenuation bias. The presence of measurement error increase the variance term in the denominator of the least square estimator and




FIGURE A4. Responses to a mismeasured technology shock

NOTE. This figure shows the distribution of the estimated responses of output, consumption, investment and hours worked to a technology shock that is subject to classical measurement error and that has been estimated using either a LP (panel A) or SVAR (panel B) approach with $p = 4$ lags. The solid line is the median response, while the dash line is the structural IRF and the dotted line is the population LP/SVAR response with infinite lags.

consequently biases the population response from the $\text{VAR}(\infty)/\text{LP}(\infty)$ towards zero. This effect is even more pronounced on the estimated IRFs in a finite sample and using finite lags, as shown by the distribution of LP- and SVAR-IRFs. Importantly, notice that attenuation bias is a problem regarding identification and hence it is common to both estimation approaches.

A.2.3.2. *The correlated government spending shock*

Figure A5 show the estimated responses to a fiscal policy shock that uses the correlated government shock, rather than the innovation, and without controlling for TFP. Hence, they can be interpreted as the responses to an identified shock that is subject to correlated measurement error and hence that breaks the exogeneity assumption. As a result, the structural IRFs (dashed lines) and the population responses (dotted lines) do not agree. Therefore, similarly to the recursive monetary policy shock and the uncorrelated proxies, the estimated IRFs are also misspecified.

Moreover, these responses coincide with those that one would have obtained by assuming that government spending is exogenous and therefore does not affect other endogenous variables contemporaneously. In other words, the recursively identified




FIGURE A5. Responses to a measured fiscal policy shock

NOTE. This figure shows the distribution of the estimated responses of output, consumption, investment and hours worked to a fiscal policy shock that have been estimated using either a LP (panel A) or SVAR (panel B) approach and $p = 4$ lags. The solid line is the median response, while the dash line is the structural IRF. Note that these responses are identical to the recursive identified fiscal policy shock that orders government spending first in the VAR.

government spending shock leads to the same IRFs as the external proxy that is correlated with technology. Hence, the estimation results from targeting these estimated responses are identical.

A.2.3.3. Unit normalization with uncorrelated external proxies

Figure A6 shows the responses of output, consumption, investment and hours worked to a 1 percentage point increase in the policy rate. These responses have been obtained after implementing the unit effect normalization of Stock and Watson (2018). That is, the size of the shock has been normalized to unity using the initial impact of the shock on the policy variable, i.e. the policy rate in the case of monetary policy.

As shown by the aforementioned figure, the unit effect normalization helps eliminate the attenuation bias incurred in the estimation that uses proxies that have uncorrelated measurement error. In fact, one sees how the structural IRFs (dashed lines) and the population responses (dotted lines) agree at all horizons and for all variables. Note that these responses are identical to those obtained when employing the true innovation of the shock – see Figure A1 – if they were rescaled by a constant factor that captures the size of the shock.




FIGURE A6. Unit normalized responses to a measured monetary policy shock

NOTE. This figure shows the distribution of the estimated responses of output, consumption, investment and hours worked to a 1% increase in the real interest rate that have been estimated using either a LP (panel A) or SVAR (panel B) approach and $p = 4$ lags. The solid line is the median response, while the dash line is the structural IRF.

The estimated responses also present the bias variance trade-off which is not affected by the normalization of the size of the shock.

Appendix A.3. Hyperparameter Choices

A.3.1. Lag Length

The number of lags used in the VAR or as controls in the LP is a fundamental choice that may shape the dynamic response to shocks. Hence, given its relevance for the targeted responses used in a minimum distance estimator, such as the ones considered in this paper, it is also crucial for understanding the structural estimated parameters and the performance of the estimation as a whole.

To shed light on this issue, I plotted the response of output to a monetary policy shock estimated by LP and SVAR models under four different choices of the lag length $p \in \{2, 4, 8, 12\}$ in Figure A7. It shows that: (i) impulse responses estimated with the observed innovation and using LPs are independent of the lag length, and consequently, the median and the confidence intervals are similar across the four panels; (ii) SVAR-IRFs approximately agree with the LP-IRFs up to horizon $h \leq p$ as shown in Plagborg-Møller and Wolf (2021); and (iii) the SVAR confidence intervals converge to those of the




FIGURE A7. Output responses to an observed monetary innovation

NOTE. This figure plots the response of output to a monetary policy shock when it is estimated using either LPs (red) or SVAR (blue) under different choices of the lag length $p \in \{2, 4, 8, 12\}$. The solid line is the median response while the dash lines are the 5th and 95th percentiles coming from the different draws of the shock.

TABLE A1. Overall performance & lag length

	IRF matching				Indirect Inference			
	J_{irf}	J^*	Time	J_{unt}^*	J_{smm}	J^*	Time	J_{unt}^*
$p = 2$								
<i>Local Projection</i>	35.75	0.24	3.30 min	18.97	25.47	0.34	18.93 min	18.02
<i>Structural VAR</i>	34.61	0.61	4.32 min	17.00	26.25	0.16	11.88 min	19.32
$p = 4$								
<i>Local Projection</i>	35.68	0.25	3.40 min	18.74	30.26	0.37	28.99 min	17.95
<i>Structural VAR</i>	36.01	0.39	3.89 min	17.75	31.49	0.26	15.35 min	18.26
$p = 8$								
<i>Local Projection</i>	34.69	0.28	3.83 min	18.47	35.91	0.44	45.06 min	17.69
<i>Structural VAR</i>	34.92	0.34	3.85 min	18.36	37.26	0.49	13.35 min	18.01
$p = 12$								
<i>Local Projection</i>	34.27	0.29	3.44 min	18.63	38.52	0.41	78.53 min	17.98
<i>Structural VAR</i>	35.39	0.30	3.67 min	18.61	40.47	0.41	17.29 min	17.98

NOTE. This table breaks down the overall performance of the two econometric models in the two estimation strategies by the lag length.

LP as suggested by the theoretical results in Olea et al. (2024). In fact, they also show that increasing the lag length ameliorates the VAR coverage, but at the cost of delivering intervals as wide as those of LP.

These properties are also present when analyzing the responses to other variables as well as other shocks. Hence, they are general enough to help us understand the role of p in the minimum distance estimation that uses IRFs as targets or data moments. Table A1 breaks down by lag length the metrics presented in Table 1.2 in the main text, where recall I was averaging across different sources of variation as well. As mentioned in Section ??, the J^* , the preferred measure of overall performance, is very similar across both econometric models and in both estimation approaches when the lag length is big enough $p = 12$. However, the explanation on why J^* gets closer between LPs and SVARs is very different depending on the DSGE estimation method. For the *IRF matching* approach, it is the reduction of the bias in the SVAR-IRFs as p gets large that reduces the value of J^* and consequently the bias of the estimated parameters; while for the *Ind. Inf.* approach it is the increase in the variance of the SVAR-IRFs as p gets large that explains the increase in J^* until it converges to the level of the J^* associated with the LP approach.

From this table, one can also learn that the estimation time is independent of the lag length in the *IRF matching* approach because the model counterpart of the targeted IRFs are the structural responses which are independent of p . However, for the *Ind. Inf.* approach, the computation time is increasing in p as it requires to estimate more coefficients in each iteration of the minimization problem. This issue is even more acute in the LP approach as its flexibility is associated in part to the larger number of estimated coefficients. Overall, these results seem to call for estimating DSGE models by *Ind. Inf.* and using a VAR with small p as the auxiliary model. Nonetheless, if computational time turns to be a problem, resorting to *IRF matching* while targeting LP-IRFs becomes the second best.

A.3.2. Sample Size

Herbst and Johannsen (2023) have shown that LP can be severely biased in small samples and proposed an approach to correct for it. I investigate the consequences of this finding, as well as their proposed solution, in the context of DSGE estimation that uses estimated IRFs as targets / data moments in a minimum distance optimization. To shed light on the issue I plot in Figure A8 the estimated output response to a technology shock using LP and SVARs as well as their bias corrected counterparts for different sample sizes. In all scenarios, the simulated sample comes from the Smets and Wouters model at the true parameter vector and the lag length is set to $p = 2$. Focusing initially on the Least Squares LP (solid red line), one sees that the smaller T is, the larger the small sample bias is, and it is only at very large T s when the estimated response follows closely the structural IRF at all horizons. Moreover, the bias correction LP model of Herbst and Johannsen (2023), depicted by the dashed orange line, partially corrects for the bias in the estimated responses and are closer to the true structural IRF at all horizons and all sample sizes, which validates their Monte Carlo results for a different DGP. Moving into the SVAR-IRFs, one sees that increasing the sample size does not decrease the higher bias relative to the LP. In fact, the SVAR-IRF is very similar across all samples. Nonetheless, the bias correction term from Pope (1990) reduces the bias of the response and brings it closer to the structural IRF.

Small sample uncertainty is not only concerning in terms of bias, but also in terms of variance. As shown in Figure A9, the fan chart that depicts the distribution of output, consumption, investment and hours worked responses to a monetary policy shock are wider relative to those in Figure A1, which were estimated on a sample with $T = 300$ observations. Hence, the lower sample size can potentially impact the outcomes of both





FIGURE A8. Small sample size & bias correction

NOTE. This figure plots the response of output to a technology shock. The black dash line is the structural IRF at the true parameter vector Θ^* . The other IRFs are estimated using $T = 300$ observations (panel A) or $T = 100$ observations (panel B) by means of Least Squares LP (solid red), Bias Corrected LP (dashed orange), Least Squares VAR (solid blue) or Bias Corrected VAR (dashed purple). To give context to the role of sample size, panels C and D also plot these IRFs for $T = 500$ and $T = 1000$, respectively.

estimation approaches considered in this paper. Intuitively, the increased bias has a larger bite in the *IRF matching* approach, while the increase uncertainty affects the *Ind. Inf.* more as this approach is robust to misspecification in the auxiliary econometric model.

The implications of these results for the overall performance of the estimation are shown in Table A2. As already discussed in Section 1.5.1, I distinguish between the role of sample size when implementing or not the bias correction in the econometric models. If bias correction is not used, i.e. least squares still being used to estimate LP and VAR coefficients, then the smaller sample worsens the performance of the estimation for both auxiliary econometric models. This can be seen by the larger J^* when comparing the 3rd and 4th row to the 1st and 2nd row in that table. Interestingly, larger bias of targeted responses affects differently *IRF matching* and *Ind. Inf.* approaches. Recall that the later is robust to misspecification. Hence, *Ind. Inf.* outperforms *IRF matching* when the small sample bias in LPs is sufficiently large as shown by the smaller J^* in the 3rd



A. Local Projections

B. SVARs

FIGURE A9. Counterpart of Figure A1 with $T = 100$ observations

NOTE. This figure is the counterpart of Figure A1 when using $T = 100$ observations, instead of $T = 300$, to estimate the impulse responses to a monetary policy shock.

row. In other words, if the DSGE modeler suspects that her targeted IRFs can suffer from small sample issues, she will be better off by estimating her model using *Ind. Inf.* techniques.

TABLE A2. Overall performance & sample size

	IRF matching				Indirect Inference			
	J_{irf}	J^*	Time	J_{unt}^*	J_{smm}	J^*	Time	J_{unt}^*
$T = 300$								
<i>Local Projection</i>	35.10	0.27	3.49 min	18.70	32.54	0.39	42.88 min	17.91
<i>Structural VAR</i>	35.23	0.41	3.93 min	17.93	33.87	0.33	14.47 min	18.39
$T = 100$								
<i>Local Projection</i>	29.71	0.53	3.56 min	18.13	22.00	0.46	18.46 min	19.03
<i>Structural VAR</i>	31.62	0.47	3.33 min	17.98	25.16	0.36	9.78 min	19.50
<i>Bias Corrected LP</i>	31.55	0.32	3.26 min	19.18	23.29	0.35	20.48 min	19.50
<i>Bias Corrected SVAR</i>	33.48	0.32	3.42 min	18.65	26.06	0.33	11.02 min	20.11

NOTE. This table show the overall performance of the estimation when using two different sample sizes to generate the data moments / targets as well as the role of bias correction terms in the estimation of IRFs and its implications for the estimation outcomes.

Regarding the use of bias correction terms such as those proposed by Herbst and Johansson (2023), the second block of Table A2 shows that they can be very useful in the context of *IRF matching*. In fact, the J^* is significantly lower when using bias corrected responses as targets. For *Ind. Inf.* bias correction seems not to be very relevant as the overall outcome, specially for VARs, is similar to that obtained without bias correction terms. This is a puzzling result as bias correction comes at the cost of higher variance, but this increase in IRF uncertainty doesn't seem to reflect on the structural parameters.

A.3.3. Weighting Matrices

The selection of the weighting matrix may have a substantial impact on the estimation outcome. Differently from the previous hyper-parameter choices, this choice affects *IRF matching* and *Ind. Inf.* in the same way. Therefore, in Table A3 I only report the mean and standard deviation of each estimated parameter via the *Ind. Inf.* approach, as the same lessons apply to the *IRF matching* estimates.

First, we have seen that J^* 's decreased as we move away from the identity matrix. This is reflected in the lower bias in key parameters for capturing the dynamic responses to the targeted shocks. In fact, the bias of the inter- and intra-temporal elasticities of substitution and the habit parameter $\{\sigma_c, \sigma_l, h_c\}$ present when using the identity matrix almost disappear when using the optimal weighting matrix. In fact, these three parameters are crucial for capturing the dynamic response to aggregate shocks in

TABLE A3. Indirect Inference Estimated Parameters

Parameter	Truth	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
		Identity Matrix		Diagonal Matrix		Optimal Matrix	
$\hat{\sigma}_c$	1.38	1.24	0.38	1.25	0.38	1.37	0.38
$\hat{h}$	0.71	0.77	0.16	0.78	0.12	0.73	0.15
$\hat{\sigma}_l$	1.83	1.88	0.59	1.89	0.59	1.84	0.57
$\hat{\phi}$	5.74	5.44	1.82	5.43	1.78	5.11	1.77
$\hat{\xi}_w$	0.70	0.62	0.20	0.62	0.20	0.62	0.19
$\hat{\xi}_p$	0.66	0.66	0.20	0.65	0.19	0.64	0.20
$\hat{\imath}_w$	0.58	0.55	0.19	0.56	0.19	0.57	0.19
$\hat{\imath}_p$	0.24	0.23	0.08	0.23	0.08	0.23	0.08

NOTE. This table depicts the true value of the estimated parameters from the Smets and Wouters model. It also displays the mean and standard deviation of each parameter under the three analyzed weighting matrices. The values of the mean and standard deviation are the average and the maximum across the different sources of variation and lag lengths considered, respectively.

the Smets and Wouters model. And second, it seems that the standard deviations of these parameters are not affected by the choice of the weighting matrix. The optimal weighting matrix slightly improves the efficiency of the estimation but not as much as initially expected.

Appendix B

Appendix to Chapter 2

Appendix B.1. Additional empirical evidence

B.1.1. Irish rental sector

In our model economy we assume that the rental sector is populated by households that own one or two rental properties. Although, this assumption may seem restrictive, it is consistent with the Irish private rental sector.

In Figure A1, we use data from the Central Statistical Office (CSO) on residential property transactions to show that the vast majority of non-occupier property purchases



FIGURE A1. Share of property transactions, by type of buyer and year

NOTE: This figure shows the share of all house sales by type of buyer and year in panel A. Panel B focus the attention in non-occupier buyers which are split into two categories: household buyers and non-household buyers. Data is available at the CSO.




FIGURE A2. Irish rental sector structure

NOTE: This figure shows the share of landlords by number of registered tenancies (panel A) as well as the share of tenancies owned by landlords with different number of registered tenancies (panel B). Data is from the RTB.

correspond to household buyers. In fact, in 2015, the year when the macro-prudential reform was introduced, around 70% of those transactions correspond to household buyers. Nonetheless, these data also confirms that the role of non-household buyers such as pension funds, private rental firms and Real Estate Investment Trusts (REITs) has increased over the last decade.

In Figure A2, we dig deeper into the ownership structure in the rental sector and use data from the Residential Tenancies Board (RTB) using RTB registrations as a proxy for ownership. Panel A shows the share of landlords by number of tenancies. Note that a tenancy is not fully analogous to a property as there may be some instances where there are multiple tenancies in one property (e.g. a flat with multiple rented rooms). Nonetheless, the vast majority of tenancies are individual properties. With that in mind, the evidence on the RTB data points to a lesser role of large scale professional landlords as only 4.6% of tenancies are held by landlords with more than 100 units. On the other hand, the vast majority of landlords register a single rental property (70%) or at most two (86%). One gets a similar picture, if looks at the share of tenancies by landlords – panel B. In fact, landlords with one or two properties registered more than 50% of all tenancies.

Figure A1 and A2 are consistent with each other as the rise of institutional investors in Ireland is mostly concentrated in newly constructed, high quality and well located units, but not so relevant at the aggregate level (Ireland's Department of Finance 2019).

B.1.2. Macro-prudential limits, house & rental prices

In this section, we describe the data used in our regression analysis, provide additional non-parametric evidence on the opposite response of house and rental prices to the introduction of macro-prudential limits, and run some robustness test that verify such relationships.

B.1.2.1. Data sources

The core of our final data set is the result of combining the “distance” measure with county-level selling and rental house prices. In our main specifications, we borrow the “distance” measure from Acharya et al. (2022). They construct this measure using loan-level information on residential mortgages. In particular, they “calculate what would have been the distance from the limits for each borrower in the year before the policy, assuming that the limits were in place during that period” (p. 12, Acharya et al., 2022). For confidential reasons, we got this information aggregated to the county level.

Data on house and rental prices comes from Daft.ie. We borrow these data from Lyons (2018) since in his website he has the aggregated time series for each Irish county of both selling and rental prices.




FIGURE A3. Counties, lending limits, house & rental price growth

NOTE. This figure shows the county-level distance from the limits (panel A), house price (panel B) and rental price (panel C) growth between the third quarter of 2014 and the fourth quarter of 2016. Data on prices comes from Daft.ie while the distance measure was provided by Mateo Crosignani and corresponds to the one in their paper: Acharya et al. (2022). Darker colors indicate less distant counties, lower house price growth and higher rental price growth.

B.1.2.2. Non-parametric evidence

Figure A3 shows the variation in house price growth (panel A), the distance measure (panel B) and rental price growth (panel C) across all Irish counties. In low-distance counties, such as those areas around Dublin, house price growth was slower and close to zero, while rental prices were growing faster at a pace around 30-35%. This observation suggests that the distance measure is positively correlated with house price growth while it is negatively correlated with rental price growth. This statement was formally verified in our regression analysis in Section 2.3.1.

Appendix B.2. Further model details

B.2.1. Solution method

The steady state solutions of the model consist of two main loops: an inner loop that solves the household problem given structural parameters and prices, and an outer loop that recovers the equilibrium distribution and prices. A description of the algorithms used for the approximation of the steady state equilibria can be found in Appendix B.2.1.1 and B.2.1.2. In addition to steady state equilibria, welfare comparisons also require to solve the transition from one steady state to another. The computational approach used to solve for such transition is described in Appendix B.2.1.3.

B.2.1.1. Household problem

As shown in Section 2.2.2, the household state variables are age, j , income, y , the housing state, $s = (h, \tilde{h})$, and net financial wealth, a . Consequently, the first step is to discretize the continuous state variables. Financial wealth lie on a non-linearly spaced grid with 150 points that includes 50 negative values and 100 positive ones, while the stochastic component of income is discretized using the approach in De Nardi, Fella, and Paz-Pardo (2020) that accounts for non-linearities and age-dependence. In particular, we allow for 7 points for the stochastic component of income whose values vary with the working age of the household.¹ The remaining state variables are already discrete. Recall that: (a) the model period is one year and household live up to 71 years, and (b) the housing state can take 5 different values: (i) $s = (0, \tilde{h}_1)$ if renter, (ii) $s = (1, \tilde{h}_1)$ if small owner, (iii) $s = (1, \tilde{h}_2)$ if big owner, (iv) $s = (2, \tilde{h}_2)$ if landlord with one rented house, or (v) $s = (3, \tilde{h}_2)$ if landlord with two rented houses.

Since households die with certainty at age J , we know their optimal policy in their terminal period, so we can proceed by backward induction and compute the remaining age-dependent policy functions. Note that households make the standard consumption-savings choice, a' , as well as decide on the next period housing tenure, s' , at each age. Given that the housing choice is discrete, the solution of the household problem requires using computational techniques employed to solve discrete-continuous dynamic choice models. We follow closely the recipe from Fella (2014) and Iskhakov et al. (2017) to use the endogenous grid method (EGM) together with taste shocks to solve for these discrete-

¹ The transition matrix that controls the evolution of household's income over time is also age-dependent and hence it is of dimension $7 \times 7 \times 41$ where 41 is the retirement age J^{ret} .

choice specific policy and value functions. In a nutshell, for each $j < J$ we first compute the expected marginal utility to then invert the Euler equation and get the endogenous consumption-asset policy in a normal EGM step. After that, we apply the general EGM procedure to verify the global optimality of these choices in the non-concave region and discard those that are not. Finally, we use the obtained s' -dependent value and policy functions to compute the probability of the discrete choice using the Logit probability formula and the expected value function using the log-sum formula. These are stored and used in the next step of the backward induction. Once the backward induction is finished, the final outcomes of the algorithm are s' -dependent consumption-savings policy functions, a discrete choice probability and a value function.

B.2.1.2. General equilibrium

To compute the equilibrium in the housing and rental markets we proceed as follows:

1. Make a guess for the rental price, p_r^g .
2. Make a guess for the share of low quality housing, $H_2^{sh,g}$. Note that this allows us to know the share of the high quality housing H_1^{sh} as both sum up to one. Recall that there are no empty houses and population size is normalized to 1.
3. Use these guessed shares to recover the transaction prices of the two qualities, $p^g(\tilde{h}_1)$ and $p^g(\tilde{h}_2)$, using the equilibrium condition (2.20). Note that the only endogenous object in that expression is the aggregate housing stock H , which only changes because of the equilibrium share of each quality type. Recall that $\tilde{h}_1$ and $\tilde{h}_2$ are fixed during calibration.
4. Given price guesses $\{p_r^g, p^g(\tilde{h}_1), p^g(\tilde{h}_2)\}$, use the algorithm loosely described in Appendix B.2.1.1 to get the value and policy functions that solve the household problem.
5. Using the household's consumption-saving policy and the discrete choice probability, recover the stationary distribution of households $\mathcal{D}(a, s, y, j)$ as it contains all the information needed for evaluating if the rental and housing market clear.
 - a. Rental demand equals the share of households that choose to be renters

$$R^d = \sum_{i_a=1}^{n_a} \sum_{i_y=1}^{n_y} \sum_{j=1}^J \mathcal{D}(a_{i_a}, s_1, y_{i_y}, j)$$

- b. Rental supply is given by the sum of landlords with one rented out property plus two times the share of landlords with two rented out properties.

$$R^s = \sum_{i_a=1}^{n_a} \sum_{i_y=1}^{n_y} \sum_{j=1}^J \mathcal{D}(a_{i_a}, s_3, y_{i_y}, j) + 2 \times \sum_{i_a=1}^{n_a} \sum_{i_y=1}^{n_y} \sum_{j=1}^J \mathcal{D}(a_{i_a}, s_5, y_{i_y}, j)$$

- c. The share of households living in the low quality home is also given by the equilibrium distribution

$$H_2^{sh,d} = \sum_{i_a=1}^{n_a} \sum_{i_y=1}^{n_y} \sum_{j=1}^J \mathcal{D}(a_{i_a}, s_1, y_{i_y}, j) + \sum_{i_a=1}^{n_a} \sum_{i_y=1}^{n_y} \sum_{j=1}^J \mathcal{D}(a_{i_a}, s_2, y_{i_y}, j)$$

6. If $|R^d - R^s| < \varepsilon_r$ and $|H_2^{sh,g} - H_2^{sh,d}| < \varepsilon_h$, then we are done. Otherwise, we need to update the guesses and go back to step 3. For the share of low quality houses, we use the convex combination between the previous guess and the solution from the household problem, while for the rental price we increase the guess if $R^s < R^d$ and decrease it otherwise.

Hence, the final outcomes of this algorithm are: an equilibrium rental price, an equilibrium average house price, the stationary distribution of households over their state space and optimal policy and value functions.

B.2.1.3. Transition dynamics

To compute the transition paths shown in Figures 2.4 and 2.7, we resort to the traditional approach that assumes that at time $t = 0$ the economy is initially in a steady state. Then, at $t = 1$ the policy reform is introduced as a surprise for households and maintained forever. Recall that in the macroprudential experiment the policy reform consists in introducing tighter LTV and LTI limits, so that λ_{LTV} and λ_{LTI} change while everything else remains untouched; while for the interest rate experiment, it is only the return on financial assets r_s and the mortgage rate r_b that increase in the new steady state. In either case, the key idea is to assume that after T periods the transition from the old to the new steady state is completed. As a result, one can safely assume that policy and value functions at time $t = T$ are those from the new steady state. So that $c_T = c_{ss}^{new}$, $a_T = a_{ss}^{new}$ and $\mathbb{P}_T(s) = \mathbb{P}_{ss}^{new}(s)$.

For a given sequence of prices $\left\{ p_t^r, p_t(\tilde{h}) \right\}_{t=1}^T$, the previous insight allow us to solve the household problem backwards and obtain their policy functions at each point in

time $\{c_t, a_t, \mathbb{P}_t(s)\}_{t=1}^T$. Knowing that $\mathcal{D}_0 = \mathcal{D}_1$, these are useful to iterate the distribution forward: $\mathcal{D}_{t+1} = \Gamma_t(\mathcal{D}_t)$ where Γ_t is the mapping obtained from the policy functions. Finally, using the sequence of household distributions over their state space $\{\mathcal{D}_t\}_{t=0}^T$ one can check if rental and housing markets clear at each point in time. If they do not, then the given sequence of prices needs to be updated until they do.

Thus, the most difficult aspect of the transition is to find suitable paths for rental and house prices. We approach this problem by first guessing different rental and housing price paths and evaluating ex-post which ones are closer to form an equilibrium sequence in housing and rental markets. These guesses are constructed parametrically by imposing an initial jump and a degree of curvature in its reversal to the new steady state level. Once we have a sense on how these equilibrium paths should look like, we follow a similar approach to that described in point 6 of the general equilibrium algorithm with the caveat that we now update the guesses based on the gaps between supply and demand along the entire path and not just based on one point in time.

B.2.2. LTI and LTV implementation in Ireland

As stated in Section 2.2.2, the borrower must satisfy two constraints. First, a loan-to-income (LTI) requirement that limits household's borrowing to a multiple, λ_{LTI} , of its current (annual) income. And second, a maximum loan-to-value (LTV) limit, which imposes that the size of the mortgage has to be smaller than a fraction of the value of the house.

When Central Banks establish these limits, they often include some exemptions based on the type of borrower or the type of property households purchase. For example, the Irish reform of 2015 imposed a LTI limit of 3.5 that only applied to First Time Buyers (FTBs). In the model, we identify FTBs with households that transition from renting into owning as there are very few (or even zero) households that after selling their primary residence become homeowners for a second time during their life-cycle. For all other borrowers, we let the pre-reform limit to apply as this was the LTI implicitly imposed by banks in absence of the Bank of Ireland macro-prudential framework. Hence, formally, the LTI in the *post-reform* economy is

$$\begin{aligned}
 \text{(A1)} \quad & a' \geq -\lambda_{LTI}^{post} y & \text{if } h' = 1 > h \\
 \text{(A2)} \quad & a' \geq -\lambda_{LTI}^{pre} y & \text{otherwise}
 \end{aligned}$$

Moreover, the reform also included some exceptions for the LTV limit based on the type of purchase. For example, buy-to-let buyers faced a more stringent 70% loan-to-value limit. We include this feature in the model by distinguishing between owner-occupied and buy to let purchases for which we let $\lambda_{LTV}^{oo} = 0.8$ and $\lambda_{LTV}^{btl} = 0.7$ to apply. In the model, it is easy to identify this purchases as households that one more than one property always lease it out. Hence, in the *post-reform* economy the LTV limit is given by

$$(A3) \quad a' \geq -\lambda_{LTV}^{oo} p(\tilde{h}') \quad \text{if } h' = 1, h = 0$$

$$(A4) \quad a' \geq -\left(\lambda_{LTV}^{oo} p(\tilde{h}') + \lambda_{LTV}^{btl} p(\tilde{h}'_1)h'\right) \quad \text{if } h' > 1 \geq h.$$

Appendix B.3. Additional model results and experiments

B.3.1. Understanding house and rental price responses

We have seen that rental prices rise and house prices fall in response to a tightening in credit conditions. Figure A4 is useful for explaining these price dynamics as it plots the cross sectional distribution of the housing state in the *pre-* and *post-reform* economies, as well as the flows in and out of these states.

At the new tighter limits, many pre-reform homeowners cannot afford to purchase a house and therefore are pushed into renting. In order to meet that additional rental demand, rental prices need to rise to incentivize some households to purchase buy-to-let properties. Motivated by the higher rental prices, many homeowners of the better quality homes transition into the landlord state. This flow, from OWNER (big) to LANDLORD 1P, compensates about 75% of the increase in the rental demand. To a lesser extent, some existing landlords also buy additional buy-to-let properties. Overall, rental prices need to increase by 2.84% to push homeowners into the landlord state and meet the extra rental demand.




FIGURE A4. Housing flows – Pre vs. Post Reform

NOTE. This figure shows the equilibrium share of households in each housing state for the pre- and post-reform economies. Flows from a house state to another as credit conditions change helps explaining the response of rental and house prices.

Average house prices in this economy fall due to a change in the composition of the aggregate housing stock. First, the increase in rental demand results in a higher share of buy-to-let properties in the post-reform equilibrium. Since these are of lower quality and cheaper, the average house price falls. Moreover, tighter credit also pushes some households into buying the low quality house instead of the better quality one since it requires a smaller mortgage. This is shown by the flow from OWNER (big) to OWNER (small). Nonetheless, downsizing is quantitatively small in this experiment and thus there is only a small increase in the share of low quality homes in the post-reform equilibrium, leading to an also small fall in the average house price.

To highlight the big role that downsizing may have in the evolution of the average house price, we repeat the simulation used to generate the flow chart above but now for the interest rate experiment. The flows across the *low* and *high interest rate* economies are shown in Figure A5. It is easy to see that now there are more households that are OWNER (big) in the old steady state (low interest rate economy) that become OWNER (small) in the new steady state (high interest rate economy). As a result, and as we have seen in section 2.4.1.1 the average house prices fall significantly more (-1.62%) despite the smaller fall in the homeownership rate (-0.92 p.p.).




FIGURE A5. Housing flows – High vs. Low Real Interest Rate Economies

NOTE. This figure shows the equilibrium share of households in each housing state for the high and low interest rate economies. Flows from a house state to another as interest rate change helps explaining the response of rental and house prices.

B.3.2. The macro-prudential reform: interaction between limits

Loan to Value (LTV) and Loan to Income (LTI) limits are often introduced jointly, as it was the case for Ireland. To understand the contribution to each of them to the overall quantity and price effects of the reform, we compute two counterfactual economies: (i) the *Only LTI* economy which imposes the institutional 3.5 LTI limit but leaves the LTV unchanged, and (ii) the *Only LTV* economy which imposes the institutional 80% LTV limit but leaves the LTI unaltered.

Table A1 shows the change relative to the *pre-reform* economy in the rent-to-price ratio as well as the homeownership rate after imposing both (full-reform) or one of these two limits. Results show that the LTI alone has a larger effect than the LTV, but these effects are smaller than those obtained after imposing the two limits jointly. Hence, there is an interaction between the two as it has been highlighted in Greenwald (2018).

TABLE A1. Non-linear interactions between credit limits

	Full-Reform	Only LTI	Only LTV
$\Delta\%$ Rent-to-Price	+2.82 %	+1.71 %	+0.73 %
Δ Homeownership rate	-1.83 p.p	-1.13 p.p.	-0.53 p.p.

NOTE. This table shows the effects of the reform on the rent to house price rate and home-ownership rate (first column) and decomposes the role of each limit by imposing one at a time. A tighter LTI (second column) has a larger effect than the tighter LTV (third column) if they are introduced on their own.

Appendix C

Appendix to Chapter 3

Appendix C.1. Further Model Details

C.1.1. Fixed & adjustable rate mortgage economies

Fixed and adjustable rate mortgages have been an object of theoretical comparison for many years given their simplicity and predominance in many countries. Hence, we use them as a benchmark to which compare the HRM analyzed in this paper. This section shows the main equilibrium conditions of the model presented in Section ?? when all mortgages are assumed to be of either type: FRM or ARM.

When considering these economies, the only difference in the model specification concerns the evolution of mortgage payments. As shown in Greenwald (2018), borrower's promised payments in ARM and FRM are defined as

$$(A1) \quad x_{b,t}^{ARM} = q_t^* m_t$$

$$(A2) \quad x_{b,t}^{FRM} = \rho q_t^* m_t^* + (1 - \rho)(1 - \nu)\pi_t^{-1} x_{b,t-1}^{FRM}$$

respectively. Note that these two equations imply that promised payments is not a state variable in the ARM economy as these change period-by-period, while in the FRM economy it is. In fact, the promised payments in the previous period $x_{b,t-1}^{FRM}$ summarizes the entire history of payments given the recursive definition in (A2). For the saver, these only differ in the mortgage rate which is assumed to be net of taxes, i.e. $(q_t^* - \Delta_{q,t})$.

These different specifications of the law of motion of promised payments impact household's optimality conditions as shown in Section 3.3.4. Here we show how to derive those conditions using the Lagrangian as a representation of the borrower's and saver's

optimization problem. Appendix C.1.2 repeats these derivations in the HRM economy starting from the simplest case with one-period fixes.

C.1.1.1. FRM economy

Borrower's problem. It can be represented by the following Lagrangian

$$\begin{aligned}
 \mathcal{L}_{b,t}^{FRM} = & u(c_{b,t}, h_{b,t-1}, n_{b,t}) + \beta_b \mathbb{E}_t V_b(m_t, h_{b,t}, x_{b,t}) + \lambda_{b,t} \left((1 - \tau_y) w_t n_{b,t} + \right. \\
 & - \pi_t^{-1} \left((1 - \tau_y) x_{b,t-1}^{FRM} + v m_{t-1} \right) + \rho \left(m_t^* - (1 - v) \pi_t^{-1} m_{t-1} \right) + \\
 & - \delta p_t^h h_{b,t-1} - \rho p_t^h \left(h_{b,t}^* - h_{b,t-1} \right) + T_{b,t} - c_{b,t} + \\
 & \left. + \mu_t \rho \left(\bar{m}_t^{PTI} \int^{\bar{e}_t} e_i d\Gamma_e(e_i) + \bar{m}_t^{LTV} (1 - \Gamma_e(\bar{e}_t)) - m_t^* \right) \right)
 \end{aligned}
 \tag{A3}$$

where $V_b(m_t, h_{b,t}, x_{b,t})$ is the next period borrower's value function and $\lambda_{b,t}$ is the Lagrange multiplier associated to the borrower's budget constraint. Hence, the first order condition with respect to newly issued mortgages, $\frac{\partial \mathcal{L}_{b,t}^{FRM}}{\partial m_t^*} = 0$, implies

$$\begin{aligned}
 \lambda_{b,t} \rho - \lambda_{b,t} \mu_t \rho = & \beta_b \mathbb{E}_t \left[\lambda_{b,t+1} \pi_{t+1}^{-1} \rho \left((1 - \tau_y) q_t^* + v + \rho (1 - v) \right) \right] + \\
 & + \beta_s^2 \mathbb{E}_{t+1} \left[\lambda_{b,t+2} \pi_{t+2}^{-1} \rho \left(\pi_{t+2}^{-1} (1 - \rho) (1 - v) \left((1 - \tau_y) q_t^* + v + \rho (1 - v) \right) \right) \right] + \dots \\
 & + \beta_s^k \mathbb{E}_{t+k} \left[\lambda_{b,t+k+1} \pi_{t+k+1}^{-1} \rho (1 - \rho)^k (1 - v)^k \left(\prod_{j=1}^k \pi_{t+j+1}^{-1} \right) \left((1 - \tau_y) q_t^* + v + \rho (1 - v) \right) \right] + \dots
 \end{aligned}
 \tag{A4}$$

which we can write recursively as follows

$$1 = \Omega_{b,t}^m + \Omega_{b,t}^{x,FRM} q_t^* + \mu_t
 \tag{A5}$$

where μ_t is the multiplier on the borrower's aggregate credit limit, and $\Omega_{b,t}^m$ and $\Omega_{b,t}^{x,FRM}$ are the marginal costs to the borrower of taking on an additional dollar of face value debt, and of promising an additional dollar of initial payments. These values are defined by

$$\Omega_{b,t}^m = \mathbb{E}_t \left[\Lambda_{b,t+1}^b \pi_{t+1}^{-1} \left(v + \rho (1 - v) + (1 - v)(1 - \rho) \Omega_{b,t+1}^m \right) \right]
 \tag{A6}$$

$$\Omega_{b,t}^{x,FRM} = \mathbb{E}_t \left[\Lambda_{b,t+1}^b \pi_{t+1}^{-1} \left((1 - \tau_y) + (1 - v)(1 - \rho) \Omega_{b,t+1}^{x,FRM} \right) \right]
 \tag{A7}$$

where we have used the definition of the borrower's stochastic discount factor

$$(A8) \quad \Lambda_{t,t+1}^b \equiv \beta_b \frac{u_{b,t+1}^c}{u_{b,t}^c} \quad \text{where} \quad u_{b,t}^c = \frac{\partial u(c_{b,t}, h_{b,t}, n_{b,t})}{\partial c_{b,t}}.$$

Saver's problem. It is characterized by the following Lagrangian

$$(A9) \quad \begin{aligned} \mathcal{L}_{s,t}^{FRM} = & u(c_{s,t}, \tilde{H}_s, n_{s,t}) + \beta_s \mathbb{E}_t V_s(m_t, \tilde{H}_s, x_{s,t}) + \lambda_{s,t} ((1 - \tau_y) w_t n_{s,t} + \\ & + \pi_t^{-1} x_{s,t-1}^{FRM} - \rho (m_t^* - (1 - \nu) \pi_t^{-1} m_{t-1}) - \delta p_t^h \tilde{H}_s + \\ & - (R_t^{-1} b_t - b_{t-1}) + \Pi_t + T_{s,t} - c_{s,t}) \end{aligned}$$

where $V_s(m_t, \tilde{H}_s, x_{s,t})$ is the next period saver's value function and $\lambda_{s,t}$ is the Lagrange multiplier associated to the saver's budget constraint. The first order condition with respect to newly issued mortgages $\frac{\partial \mathcal{L}_{s,t}^{FRM}}{\partial m_t^*} = 0$ is therefore given by

$$(A10) \quad \begin{aligned} \lambda_{s,t} \rho_t = & \beta_s \mathbb{E}_t \left[\lambda_{s,t+1} \pi_{t+1}^{-1} \rho ((q_t^* - \Delta_{q,t}) + \rho(1 - \nu)) \right] + \\ & + \beta_s^2 \mathbb{E}_{t+1} \left[\lambda_{s,t+2} \pi_{t+1}^{-1} \rho (\pi_{t+2}^{-1} (1 - \rho) (1 - \nu) ((q_t^* - \Delta_{q,t}) + \rho(1 - \nu))) \right] + \dots \\ & + \beta_s^k \mathbb{E}_{t+k} \left[\lambda_{s,t+k+1} \pi_{t+1}^{-1} \rho (1 - \rho)^k (1 - \nu)^k \left(\prod_{j=1}^k \pi_{t+j+1}^{-1} \right) ((q_t^* - \Delta_{q,t}) + \rho(1 - \nu)) \right] + \dots \end{aligned}$$

which we can write recursively as follows

$$(A11) \quad 1 = \Omega_{s,t}^m + \Omega_{s,t}^{x,FRM} (q_t^* - \Delta_{q,t})$$

where $\Omega_{s,t}^m$ and $\Omega_{s,t}^{x,FRM}$ are the marginal continuation benefits to the saver of an additional unit of face value and an additional dollar of promised initial payments, respectively. These values are defined by

$$(A12) \quad \Omega_{s,t}^m = \mathbb{E}_t \left[\Lambda_{t,t+1}^s \pi_{t+1}^{-1} (\rho(1 - \nu) + (1 - \rho)(1 - \nu) \Omega_{s,t+1}^m) \right]$$

$$(A13) \quad \Omega_{s,t}^{x,FRM} = \mathbb{E}_t \left[\Lambda_{t,t+1}^s \pi_{t+1}^{-1} (1 + (1 - \rho)(1 - \nu) \Omega_{s,t+1}^{x,FRM}) \right]$$

where as for the borrower we have used the definition of the saver's stochastic discount factor.

C.1.1.2. ARM economy

The borrower's and the saver's optimization problems in the ARM economy are characterized by almost identical Lagrangians. Recall that the only difference is in the definition of the low of motion of promised payments. Hence, in this economy the Lagrangian of the borrower is given by equation (A3) after substituting $x_{s,t}^{FRM}$ by $x_{s,t}^{ARM}$. As a result of this change, the first order condition with respect to newly issued mortgages now is given by

$$(A14) \quad \begin{aligned} \lambda_{b,t}\rho - \lambda_{b,t}\mu_t\rho = & \beta_s \mathbb{E}_t \left[\lambda_{b,t+1}\pi_{t+1}^{-1}\rho \left((1-\tau_y) q_t^* + \nu + \rho(1-\nu) \right) \right] + \\ & + \beta_s^2 \mathbb{E}_{t+1} \left[\lambda_{b,t+2}\pi_{t+1}^{-1}\rho \left(\pi_{t+2}^{-1}(1-\rho)(1-\nu)(\nu + \rho(1-\nu)) \right) \right] + \dots \\ & + \beta_s^k \mathbb{E}_{t+k} \left[\lambda_{b,t+k+1}\pi_{t+1}^{-1}\rho (1-\rho)^k (1-\nu)^k \left(\prod_{j=1}^k \pi_{t+j+1}^{-1} \right) (\nu + \rho(1-\nu)) \right] + \dots \end{aligned}$$

which in recursive form can be rewritten as follows

$$(A15) \quad \Omega_{b,t}^{ARM} = 1 - \mu_t$$

where $\Omega_{b,t}^{ARM}$ is the total continuation cost of an additional unit of debt in the ARM economy and is now defined as

$$(A16) \quad \Omega_{b,t}^{ARM} = \mathbb{E}_t \left[\Lambda_{t,t+1}^b \pi_{t+1}^{-1} \left((1-\tau_y) q_t^* + \nu + \rho(1-\nu) + (1-\nu)(1-\rho) \Omega_{b,t+1}^{ARM} \right) \right].$$

Similarly, the saver's optimization problem is also characterized by the Lagrangian in (A9) after substituting $x_{s,t}^{FRM}$ by $x_{s,t}^{ARM}$. Consequently, the first order condition with respect to newly issued mortgages is

$$(A17) \quad \begin{aligned} \lambda_{s,t}\rho = & \beta_s \mathbb{E}_t \left[\lambda_{s,t+1}\pi_{t+1}^{-1}\rho \left((q_t^* - \Delta_{q,t}) + \rho(1-\nu) \right) \right] + \\ & + \beta_s^2 \mathbb{E}_{t+1} \left[\lambda_{s,t+2}\pi_{t+1}^{-1}\rho \left(\pi_{t+2}^{-1}(1-\rho)(1-\nu)(1+\rho(1-\nu)) \right) \right] + \dots \\ & + \beta_s^k \mathbb{E}_{t+k} \left[\lambda_{s,t+k+1}\pi_{t+1}^{-1}\rho (1-\rho)^k (1-\nu)^k \left(\prod_{j=1}^k \pi_{t+j+1}^{-1} \right) (1+\rho(1-\nu)) \right] + \dots \end{aligned}$$

which can be rewritten in recursive form as

$$(A18) \quad \Omega_{s,t}^{ARM} = 1$$

where $\Omega_{s,t}^{ARM}$ is again the total continuation benefit of an additional unit of debt in the

ARM economy and it is given by

$$(A19) \quad \Omega_{s,t}^{ARM} = \mathbb{E}_t \left[\Lambda_{t,t+1}^s \pi_{t+1}^{-1} \left((q_t^* - \Delta_{q,t}) + \rho(1-\nu) + (1-\nu)(1-\rho)\Omega_{s,t+1}^{ARM} \right) \right] .$$

C.1.2. The hybrid rate mortgage economy

C.1.2.1. The simplest example: a one-period HRM

Here we show the step-by-step computation of the first order condition of the borrower's and the saver's optimization problem with respect to the newly issued mortgages in an economy with HRM and a fixation period of one year. It helps building the intuition for the more general case with T periods in the fixed part of the contract.

Recall that the only difference between our economies is on the mortgage payment schedule. Hence, it is useful to start writing down the low of motion of promised payments for a HRM economy with one-period in the fixed rate:

$$(A20) \quad x_{b,t}^{HRM(T1)} = \rho q_t^* m_t^* + (1-\rho)(1-\nu)\pi_t^{-1} q_{t-1}^* m_{t-1}$$

where the first summand refers to the payments made when refinancing today and therefore subject to the current mortgage rate, while the second term corresponds payments made when having refinanced the period before and consequently are fixed to the rate prevailing in the previous period.

Borrower's Problem. It be represented by the Lagrangian in (A3) after substituting x_t^{FRM} for $x_t^{HRM(T1)}$. Consequently, the first order condition with respect to newly issued mortgages in this economy is

$$(A21) \quad \begin{aligned} \lambda_{b,t}\rho - \lambda_{b,t}\mu_t\rho &= \beta_s \mathbb{E}_t \left[\lambda_{b,t+1}\pi_{t+1}^{-1}\rho \left((1-\tau_y) q_t^* + \nu + \rho(1-\nu) \right) \right] + \\ &+ \beta_s^2 \mathbb{E}_{t+1} \left[\lambda_{b,t+2}\pi_{t+1}^{-1}\rho \left(\pi_{t+2}^{-1}(1-\rho)(1-\nu) \left((1-\tau_y) q_t^* + \nu + \rho(1-\nu) \right) \right) \right] + \\ &+ \beta_s^3 \mathbb{E}_{t+2} \left[\lambda_{b,t+3}\pi_{t+1}^{-1}\rho \left(\pi_{t+2}^{-1}\pi_{t+3}^{-1}(1-\rho)^2(1-\nu)^2 \right) (\nu + \rho(1-\nu)) \right] + \dots \\ &+ \beta_s^k \mathbb{E}_{t+k} \left[\lambda_{b,t+k+1}\pi_{t+1}^{-1}\rho (1-\rho)^k (1-\nu)^k \left(\prod_{j=1}^k \pi_{t+j+1}^{-1} \right) (\nu + \rho(1-\nu)) \right] + \dots \end{aligned}$$

which we can be rewritten as follows

$$(A22) \quad 1 = \Omega_{b,t}^m + \Omega_{b,t}^{x,HRM} q_t^* + \mu_t$$

where μ_t is the multiplier on borrower's aggregate credit limit, $\Omega_{b,t}^m$ is the marginal cost to the borrower of taking an additional dollar of face value debt, and $\Omega_{b,t}^{x,HRM}$ is the marginal cost to the borrower of promising an additional dollar of initial payments and it is given by:

$$(A23) \quad \Omega_{b,t}^{x,HRM} = \mathbb{E}_t \left[\Lambda_{t,t+1}^b \pi_{t+1}^{-1} (1 - \tau_y) \right] + \\ + \mathbb{E}_t \left[\Lambda_{t,t+1}^b \Lambda_{t+1,t+2}^b \pi_{t+1}^{-1} \pi_{t+2}^{-1} (1 - \rho)(1 - \nu)(1 - \tau_y) \right] .$$

Intuitively, the marginal continuation cost to the borrower of an additional dollar of payments is a finite sum truncated at the end of the fixed period.

Saver's Problem. As for the borrower, the saver's problem is defined by the Lagrangian in (A9) after substituting x_t^{FRM} for $x_t^{HRM(T1)}$. As a result, the first order condition with respect to newly issued mortgages is given by

$$(A24) \quad \lambda_{s,t} \rho = \beta_s \mathbb{E}_t \left[\lambda_{s,t+1} \pi_{t+1}^{-1} \rho \left((q_t^* - \Delta_{q,t}) + \rho(1 - \nu) \right) \right] + \\ + \beta_s^2 \mathbb{E}_{t+1} \left[\lambda_{s,t+2} \pi_{t+1}^{-1} \rho \left(\pi_{t+2}^{-1} (1 - \rho)(1 - \nu) \left((q_t^* - \Delta_{q,t}) + \rho(1 - \nu) \right) \right) \right] + \\ + \beta_s^3 \mathbb{E}_{t+1} \left[\lambda_{s,t+3} \pi_{t+1}^{-1} \rho \left(\pi_{t+2}^{-1} \pi_{t+3}^{-1} (1 - \rho)^2 (1 - \nu)^2 (1 + \rho(1 - \nu)) \right) \right] + \dots \\ + \beta_s^k \mathbb{E}_{t+k} \left[\lambda_{s,t+k+1} \pi_{t+1}^{-1} \rho (1 - \rho)^k (1 - \nu)^k \left(\prod_{j=1}^k \pi_{t+j+1}^{-1} \right) (1 + \rho(1 - \nu)) \right] + \dots$$

which can be rewritten in the following form

$$(A25) \quad 1 = \Omega_{s,t}^m + \Omega_{s,t}^{x,HRM(T1)} (q_t^* - \Delta_{q,t})$$

where $\Omega_{s,t}^m$ is the marginal continuation benefit of an additional unit of face value debt as defined in equation (3.26) and $\Omega_{s,t}^{x,HRM(T1)}$ is the marginal benefit to the saver of promising an additional dollar of initial payments and it is defined as

$$(A26) \quad \Omega_{s,t}^{x,HRM(T1)} = \mathbb{E}_t \left[\Lambda_{t,t+1}^s \pi_{t+1}^{-1} \right] + \mathbb{E}_t \left[\Lambda_{t,t+1}^s \Lambda_{t+1,t+2}^s \pi_{t+1}^{-1} \pi_{t+2}^{-1} (1 - \rho)(1 - \nu) \right] .$$

C.1.2.2. A general T-periods HRM

Using the intuition from the simplest case we show the derivations required to obtain the first order conditions shown in Section 3.3.4 for the general HRM economy. We first start with the first order condition of the borrower's problem with respect to newly

issued mortgages

$$\begin{aligned}
\lambda_{b,t}\rho - \lambda_{b,t}\mu_t\rho &= \beta_s \mathbb{E}_t \left[\lambda_{b,t+1} \pi_{t+1}^{-1} \rho \left((1-\tau_y) q_t^* + \nu + \rho(1-\nu) \right) \right] + \dots \\
&+ \beta_s^2 \mathbb{E}_{t+1} \left[\lambda_{b,t+2} \pi_{t+1}^{-1} \rho \left(\pi_{t+2}^{-1} (1-\rho) (1-\nu) \left((1-\tau_y) q_t^* + \nu + \rho(1-\nu) \right) \right) \right] + \dots \\
(A27) \quad &+ \beta_s^T \mathbb{E}_{t+T} \left[\lambda_{b,t+T+1} \pi_{t+1}^{-1} \rho (1-\rho)^T (1-\nu)^T \left(\prod_{j=1}^T \pi_{t+j+1}^{-1} \right) \left((1-\tau_y) q_t^* + \nu + \rho(1-\nu) \right) \right] + \dots \\
&+ \beta_s^k \mathbb{E}_{t+k} \left[\lambda_{b,t+k+1} \pi_{t+1}^{-1} \rho (1-\rho)^k (1-\nu)^k \left(\prod_{j=1}^k \pi_{t+j+1}^{-1} \right) (\nu + \rho(1-\nu)) \right] + \dots
\end{aligned}$$

where here $T < k$. This condition can be rearranged and rewritten as equation (3.24). Note that, intuitively, the term $(1-\tau_y) q_t^*$ only shows up until period T as the history of payments only matter for the marginal continuation value unit then. In contrast, this term only shows up for the first summand in the ARM economy, while it always shows up for the FRM economy.

Similarly, the first order condition of the saver's optimization problem with respect to newly issued mortgages is given by

$$\begin{aligned}
\lambda_{s,t}\rho &= \beta_s \mathbb{E}_t \left[\lambda_{s,t+1} \pi_{t+1}^{-1} \rho \left((q_t^* - \Delta_{q,t}) + \rho(1-\nu) \right) \right] + \\
&+ \beta_s^2 \mathbb{E}_{t+1} \left[\lambda_{s,t+2} \pi_{t+1}^{-1} \rho \left(\pi_{t+2}^{-1} (1-\rho) (1-\nu) \left((q_t^* - \Delta_{q,t}) + \rho(1-\nu) \right) \right) \right] + \\
(A28) \quad &+ \beta_s^T \mathbb{E}_{t+T} \left[\lambda_{s,t+T+1} \pi_{t+1}^{-1} \rho (1-\rho)^T (1-\nu)^T \left(\prod_{j=1}^T \pi_{t+j+1}^{-1} \right) \left((q_t^* - \Delta_{q,t}) + \rho(1-\nu) \right) \right] + \dots \\
&+ \beta_s^k \mathbb{E}_{t+k} \left[\lambda_{s,t+k+1} \pi_{t+1}^{-1} \rho (1-\rho)^k (1-\nu)^k \left(\prod_{j=1}^k \pi_{t+j+1}^{-1} \right) (1 + \rho(1-\nu)) \right] + \dots
\end{aligned}$$

where $T < k$ and again can be rearranged and rewritten as equation (3.27).

Appendix C.2. Additional Figures

C.2.1. Temporary monetary policy shock




FIGURE A1. Response to a 1% (temporary) monetary policy shock

NOTE. Responses are normalized such that R_t increases by 1% upon impact in the HRM, FRM & ARM economies. A value of 1 represents a 1% increase relative to the steady state except for F^{LTV} , New Loan LTV, and New Loan PTI, which are measured in percentage points, and New Issuance, which is measured as a fraction of steady state output. Variable definitions are as follows. Price-to-Rent Ratio: $p_t^h/(u_t^h/u_t^c)$, Mortgage Rate: $q_t^* - \nu$, Avg. Debt Limit: $\bar{m}_t$, Debt: m_t , Mortgage payments: $\pi_t^{-1}x_{t-1}$, New Issuance: $\rho(m_t^* - (1 - \nu)\pi_t^{-1}m_{t-1})$, New Loan LTV: $m_t^*/p_t^h h_{b,t}^*$, New Loan PTI: $q_t^* m_t^*/w_t n_{b,t}$. Avg. Debt Limit, Debt, Output, Borr. Cons. and Saver Cons. are reported in real terms. Mortgage Rate, R_t , Output and Inflation are annualized.

C.2.2. Persistent inflation target shock






FIGURE A2. Response to a 1% (persistent) inflation target shock

NOTE. Responses are normalized such that R_t increases by 1% upon impact in the HRM, FRM & ARM economies. A value of 1 represents a 1% increase relative to the steady state except for F^{LTV} , New Loan LTV, and New Loan PTI, which are measured in percentage points, and New Issuance, which is measured as a fraction of steady state output. Variable definitions are as follows. Price-to-Rent Ratio: $p_t^h/(u_t^h/u_t^c)$, Mortgage Rate: $q_t^* - v$, Avg. Debt Limit: $\bar{m}_t$, Debt: m_t , Mortgage payments: $\pi_t^{-1}x_{t-1}$, New Issuance: $\rho(m_t^* - (1-v)\pi_t^{-1}m_{t-1})$, New Loan LTV: $m_t^*/p_t^h h_{b,t}^*$, New Loan PTI: $q_t^* m_t^*/w_t n_{b,t}$. Avg. Debt Limit, Debt, Output, Borr. Cons. and Saver Cons. are reported in real terms. Mortgage Rate, R_t , Output and Inflation are annualized.

C.2.3. Alternative PTI & LTV calibrations



A. Temporary Monetary Policy Shock




B. Inflation Target Shock


FIGURE A3. Loose Credit Limits & the Fixed Rate Mortgage Economy

NOTE. A value of 1 represents a 1% increase relative to the steady state, except for F^{LTV} and new issuance which are expressed in percentage points. New debt issuance is defined as: $\rho(m_t^* - (1 - \nu)\pi_t^{-1}m_{t-1})$, the house price is p_t^h , and the average debt limit $\bar{m}_t$. Output and consumption are reported in real terms.

C.2.4. Constraint Switching Effect



A. Fixed Rate Mortgage Economy



B. Adjustable Rate Mortgage Economy

FIGURE A4. Constraint Switching & Persistent Inflation Target Shocks

NOTE. A value of 1 represents a 1% increase relative to the steady state, except for F^{LTV} and new issuance which are expressed in percentage points. New debt issuance is defined as: $\rho(m_t^* - (1 - \nu)\pi_t^{-1}m_{t-1})$, the average debt limit: $\bar{m}_t$, and the house price: p_t^h .