

Borrower-Based Macroprudential Policy over the Housing Cycle

Eugenio Renedo*

CEMFI

This version: October 27, 2026

House prices shape household consumption, particularly when housing is illiquid and households are highly indebted. This paper examines whether borrower-based macroprudential policies can mitigate credit-fuelled housing booms and busts and improve welfare. I build an equilibrium macro-housing model with heterogeneous households, a frictional housing market, and long-term mortgages, and calibrate it to Spain's housing boom and bust of 1998–2014. The model shows that caps on loan-to-value, debt-to-income and payment-to-income ratios substantially dampen the bust in house prices and consumption, raising welfare at the cost of a smaller boom. These gains arise partly because lower leverage limits defaults and forced sales, mitigating the price collapse and its spillovers to other households. The instruments are not interchangeable: a debt-to-income cap delivers the largest welfare gain in the Spanish baseline, while combining instruments performs best across alternative boom scenarios. The results show that the effectiveness of borrower-based policies depends on how their borrowing limits respond to the forces driving the housing cycle.

Keywords: housing, consumption, liquidity, macroprudential policy, borrower-based measures, Great Recession.

JEL codes: E21, E44, G21, G51, R21.

* CEMFI, Casado del Alisal 5, 28014 Madrid, Spain. Email: eugenio.renedo@cemfi.edu.es. Website: erenedo.github.io. I thank Jorge Abad, Sebastián Fanelli, Nezih Guner, Josep Pijoan-Mas, Kurt Mitman, Rafael Repullo and Javier Suarez, as well as participants at the BdE Financial Stability Seminar and the CEMFI Macroeconomics Workshop for valuable comments and discussions. I gratefully acknowledge financial support from grants PREP2022-000905 (funded by MICIU/AEI/10.13039/501100011033 and ESF+), MDM-2016-0684, and CEX2020-001104-M (both funded by MICIU/AEI/10.13039/501100011033). All errors are my own.

1. Introduction

There is widespread consensus that the housing market played a central role in the Great Recession. A prolonged housing boom gave way to a sharp decline in house prices, accompanied by falling consumption, rising unemployment and declining output. Empirical evidence shows that consumption responds strongly to house prices (Mian, Rao, and Sufi 2013; Mian and Sufi 2011, 2014; Guren et al. 2021), with particularly pronounced effects among indebted and credit-constrained homeowners (Aladangady 2017; Cloyne, Ferreira, and Surico 2020; Aruoba, Elul, and Kalemli-Özcan 2022).

A common interpretation of the Global Financial Crisis is that credit expansion drove housing demand beyond the capacity of supply to respond, fuelling a sustained rise in house prices. When credit turned, prices collapsed, triggering a contraction in consumption through negative wealth effects, reduced borrowing capacity as collateral values fell, and the difficulty of selling housing assets in a falling market. In response, macroprudential authorities across advanced economies introduced a range of instruments to lean against such cycles (International Monetary Fund 2014; European Systemic Risk Board 2019; Arena et al. 2020). This paper examines whether borrower-based macroprudential policies, which restrict mortgage lending at origination by capping loan-to-value, debt-to-income and payment-to-income ratios, can curb credit-fuelled housing booms and mitigate the subsequent busts, and whether doing so improves welfare.

To answer this question, I build an equilibrium macro-housing model with heterogeneous households who buy and sell houses in a frictional market and finance their purchases with long-term mortgages. I calibrate the model to the Spanish housing boom and bust of 1998–2014. Spain provides a particularly suitable setting for this analysis for three reasons.

First, the magnitude of the housing cycle was exceptional by euro-area standards. Real house prices rose 123.2% between the end of 1997 and the end of 2007 before falling 42.9% to their trough. Real consumption rose 42.6% and subsequently fell 13.3%, while new construction more than doubled before collapsing.

Second, the boom was fuelled by credit. Convergence towards monetary union reduced the real mortgage rate from an average of 4.06% over 1996–1997 to an average of 0.61% during the boom. The median mortgage at origination required no down payment, and household mortgage debt rose from 0.23 to 0.62 of GDP.

Third, Spain is predominantly a nation of homeowners, and the homeownership rate remained broadly stable throughout the cycle. The main adjustment therefore occurred not in homeownership itself, but in the value of housing assets and the debt households carried against them.

The model captures the mortgage contracts and legal framework that prevailed in Spain during this period. More than 90 per cent of Spanish mortgages carried adjustable interest rates, and lenders

retained recourse to borrowers' income after repossession. Consequently, households that lost their homes during the bust remained liable for the part of the debt that the sale of the house did not cover. Moreover, Spain had no borrower-based macroprudential regulation in place, providing a clean benchmark against which to evaluate the effects of mortgage lending restrictions.

The calibrated model reproduces the main features of the Spanish housing cycle. Given the observed paths of interest rates, labour income and lending standards, real house prices rise 113.8% during the boom, compared with 123.2% in the data, and fall 33.7% during the bust, compared with 42.9%. Real consumption rises 38.5% against 42.6% in the data and subsequently falls 10.8% against 13.3%. Falling interest rates are the main driver of the house price boom: in isolation, they raise house prices by 72.0%, compared with 30.0% for the rise in labour income.

The subsequent collapse in consumption operates largely through the housing market. Frictions in that market make the time required to sell a house and the discount needed to complete a sale vary with macroeconomic conditions. Expected time on the market rises from 4.8 months before the boom to 11.6 months at the trough. As a result, households hit by adverse income shocks during the downturn cannot rely on selling their homes quickly to obtain cash. A decomposition that isolates each channel in turn attributes 47% of the fall in consumption to housing illiquidity, 34% to tighter collateral constraints, and 19% to the loss of housing wealth.

Borrower-based macroprudential policies would have substantially dampened the housing cycle. To quantify their effects, I introduce each instrument at the start of the boom, hold it in place throughout the cycle, and compare the resulting outcomes with the benchmark economy under the same shocks. I set each instrument at the level that maximises the welfare gains for households alive at the onset of the boom. Under this criterion, a debt-to-income cap of 3.5 times annual earnings reduces the fall in house prices from 33.7% to 20.5% and the fall in consumption from 10.8% to 5.4%. These gains come at the cost of a smaller consumption boom: consumption rises by 35.0% rather than 38.5%. Overall, the policy delivers a welfare gain equivalent to 2.06% of lifetime consumption for the average household.

The policy dampens the bust much more than the boom because its primary effect is to limit leverage. It reduces the rise in house prices by about eight percentage points, but leaves households with substantially less debt relative to the value of their homes when prices fall. This lower leverage, in turn, determines how many households are forced to sell, how many default, and how much income they lose in the years that follow. Under full recourse, the consequences of default extend beyond the loss of the house: households remain liable for the part of the debt that the sale of the house does not cover, prolonging the effects of the bust on their finances.

The three instruments are not interchangeable. At their respective welfare-maximising levels, a payment-to-income cap delivers a welfare gain of 1.85%, a loan-to-value cap 1.69%, and a debt-to-income cap 2.06%. Although the three gains differ by less than half a percentage point, and all three dampen the bust much more than the boom, the debt-to-income cap ranks first however the com-

parison is made. The key difference lies in how each instrument links permitted borrowing to its denominator. Because each cap fixes a ratio rather than the level of borrowing itself, the amount a household can borrow still varies with the underlying variable. During the Spanish boom, these variables evolved very differently.

The value of a given house rose by a factor of 2.2. A loan-to-value cap therefore allowed borrowing to grow with the housing bubble: buyers still financed the same share of the purchase price, but that price had more than doubled, leaving them with larger debts to carry into the bust. The loan that a given share of income could service at the prevailing interest rate rose by a factor of 1.4 by the end of 2007, and by 1.7 at the 2005 rate trough, as incomes rose and, for a while, rates fell. A payment-to-income cap tied to the current rate therefore loosened as rates fell and tightened as they rose. Under adjustable-rate contracts, a loan sized at the 2005 interest rate was consequently repriced at the 2008 rate. By contrast, labour income rose by a factor of 1.3, so a debt-to-income cap changed little over the cycle.

General equilibrium effects narrow these differences but do not eliminate them. By excluding some buyers, a binding cap reduces house prices and thereby limits the increase in the collateral ceiling itself. Accordingly, the rise in the loan-to-value denominator falls from 2.2 to 1.92 under the optimal loan-to-value cap. Yet the differences across instruments remain. A payment-to-income cap assessed at a stressed interest rate rather than the prevailing one, as supervisors now require, behaves more like a debt-to-income cap and delivers a welfare gain of 1.97%. The weakness of the payment-based instrument therefore lies not in its structure per se, but in indexing the permitted payment to the current interest rate.

Borrower-based macroprudential policies improve welfare because they partially correct an externality that operates through house prices. When a household takes on a larger mortgage, it marginally increases the number of households that may be forced to sell during a bust, amplifying the decline in house prices. This decline can, in turn, trap other indebted households whose homes are no longer worth enough to cover their outstanding mortgages, preventing them from selling and increasing the likelihood of default. Individual borrowers do not internalise the harm their borrowing imposes on others, and lenders do not price this external cost. As a result, households collectively take on more debt than is socially optimal. A borrowing cap addresses this externality by restricting leverage across households simultaneously, achieving a reduction in borrowing that no household would choose on its own.

The model quantifies the contribution of this externality to the welfare gains from a debt-to-income cap. If the boom were permanent, the cap would reduce welfare by 0.58% of lifetime consumption for the average household, reflecting the cost of excluding some buyers from the housing market. When the bust is introduced but house prices are held to their benchmark path, the cap instead delivers a welfare gain of 0.62%. This gain reflects the protection afforded to constrained households, worth 1.20%, net of the cost of exclusion. In general equilibrium, the welfare gain rises to 2.06%. The difference arises

from the smaller house price collapse, which benefits households beyond those directly constrained by the policy. This general equilibrium effect accounts for 70% of the total welfare gain.

The welfare gains are unevenly distributed because a policy that moderates house prices redistributes wealth from households that own a home when the boom begins towards those that buy later. Renters gain 5.34%, as they can enter the housing market at lower prices and are less exposed to the wave of foreclosures. Homeowners with a mortgage gain 3.97%, because the policy limits further borrowing during the boom and protects them from the subsequent price collapse. By contrast, outright homeowners, who account for 53.9% of Spanish households, lose 0.09%. They hold the asset whose value the policy restrains but bear none of the borrowing risks that the policy mitigates. More broadly, welfare gains decline with income and increase with leverage. Nevertheless, a modest majority of households benefits under each instrument, and the instruments rank in the same order by the share of households that gain as they do by average welfare gains.

In practice, borrower-based macroprudential policies are often applied in combination. Adding a loan-to-value cap to the debt-to-income cap raises the welfare gain from 2.06% to 2.21%, although the improvement is modest. In a boom driven by falling interest rates, the debt-to-income cap binds for almost all borrowers, leaving the loan-to-value cap with little additional work to do.

The value of combining instruments becomes clearer when the source of the boom changes. When falling interest rates drive the expansion, the ranking mirrors the Spanish experience: the debt-to-income cap delivers the largest welfare gain, while the loan-to-value cap delivers the smallest. When rising labour income drives the boom instead, the debt-to-income ceiling increases alongside incomes, while house prices rise by less. The welfare gains from the three individual instruments then converge to within a few tenths of a percentage point, and the combination delivers the largest gain. The combination is therefore the best-performing rule across both types of boom. This makes it a practical choice for macroprudential authorities, which must design policies without knowing in advance what will drive the next housing cycle.

Two features of the model economy account for much of these results: housing market frictions and the supply response. First, housing liquidity is central to both the severity of the bust and the welfare gains from regulation. Replacing the frictional housing market with a Walrasian one reduces the fall in consumption from 10.8% to 4.9% and the welfare gain from the debt-to-income cap from 2.06% to 0.49%. These results suggest that roughly half of the consumption collapse, and much of the welfare value of regulation, reflect the cost of being unable to sell a housing asset when liquidity is needed most.

Second, the response of housing supply shapes the magnitude of the subsequent price correction. New construction more than doubles during the boom in the model, mirroring the increase in residential investment observed in Spain. The additional housing stock accumulated during the boom contributes to the subsequent fall in prices, while the debt-to-income cap reduces the increase in construction by

roughly one-third.

The model takes labour income directly from the data and therefore does not capture the employment effects of the construction cycle. In particular, it omits both the jobs created during the boom and those lost during the bust. A borrowing cap would have forgone some of the employment gains during the expansion while mitigating some of the subsequent job losses. Since income losses during a downturn are more costly to liquidity-constrained households than equivalent income gains during a boom, the welfare gains reported here are likely to understate the full benefits of regulation.

Related literature. This paper contributes to two strands of literature. The first studies housing booms and busts in quantitative models with heterogeneous households. Explanations of the U.S. housing cycle emphasise the role of beliefs (Kaplan, Mitman, and Violante 2020), financial liberalisation and risk sharing (Favilukis, Ludvigson, and Van Nieuwerburgh 2017), the expansion of high-leverage lending (Corbae and Quintin 2015), the maturity structure of mortgage debt (Garriga, Manuelli, and Peralta-Alva 2019), and the relaxation of lending constraints (Justiniano, Primiceri, and Tambalotti 2019). The housing market in this paper builds on Hedlund (2016) and Garriga and Hedlund (2020), who show that endogenous housing illiquidity amplifies the transmission of house prices to consumption. It extends their framework by replacing the U.S. mortgage contract with the adjustable-rate, full-recourse mortgage characteristic of Spanish law and allowing housing supply to respond elastically. For Spain, Ferreira, Gálvez, and Pidkuyko (2024) estimate a life-cycle model of tenure, debt and consumption, attributing much of the decline in homeownership among young households during the bust to tighter credit conditions and changing income dynamics. Carro (2022), in turn, uses an agent-based model of the Spanish housing market to calibrate loan-to-value and loan-to-income limits that would have reduced the amplitude of the Spanish cycle to that observed in the United Kingdom. Relative to these contributions, this paper evaluates three borrower-based instruments within a common equilibrium framework over the observed housing cycle, with forward-looking households. It identifies and quantifies the inefficiency that these policies address and shows that housing illiquidity accounts for much of their welfare value.

The second strand evaluates borrower-based macroprudential policy. Cross-country studies find that loan-to-value and income-based limits moderate credit and house-price growth (Cerutti, Claessens, and Laeven 2017; Kuttner and Shim 2016; Akinici and Olmstead-Rumsey 2018). For Spain, Bover, Torrado, and Villanueva (2019) and Galán and Lamas (2019) document how loan-to-value, loan-to-income and debt-service ratios at mortgage origination evolved over the housing cycle, while García-Montalvo and Raya (2018) caution against relying on loan-to-value limits alone. These studies estimate average policy effects or document changes in lending distributions, but do not establish how alternative instruments would have performed over a specific historical cycle. The structural approach adopted here addresses this question directly.

Other structural evaluations include agent-based models of loan-to-value and loan-to-income limits

(Lalotitis et al. 2020; Carro et al. 2022) and models examining the interaction between collateral and payment constraints (Greenwald 2018; Grodecka 2020). Greenwald (2018) shows that payment-to-income limits govern the transmission of interest-rate cuts to mortgage credit. The corresponding implication in this paper is that this same mechanism can weaken such limits during a boom driven by falling interest rates. Finally, the inefficiency quantified here relates to the pecuniary externalities studied by Lorenzoni (2008), Bianchi (2011) and Dávila and Korinek (2018). In the present model, the externality operates through house prices, which enter households' decisions to list their homes for sale and to default as well as their collateral constraint on borrowing.

Outline. Section 2 describes the Spanish housing boom and bust of 1998–2014. Section 3 presents the model and Section 4 shows how it is taken to the data. Section 5 reproduces the boom and bust and decomposes it into its drivers. Section 6 evaluates the three instruments, decomposes the welfare gain, reports who gains, and examines combinations. Section 7 isolates the role of housing liquidity and supply. Section 8 concludes.

2. The Spanish Boom and Bust


Spain's housing cycle of 1998–2014 was among the largest any advanced economy has experienced. This section presents its anatomy: the rise and fall of house prices, consumption, homeownership, household debt and construction, and the changes in interest rates, lending standards and labour income that accompanied them. Throughout, the pre-boom reference point is the last quarter of 1997, the boom runs to the last quarter of 2007, and the bust from the first quarter of 2008 to the trough; Appendix A gives the data sources.

2.1. House Prices, Consumption, Homeownership and Construction

Real house prices in Spain and in the euro area both rose and then fell between the late 1990s and the early 2010s, but the Spanish amplitude is of a different order: real prices in Spain rose 123.2% between 1997Q4 and 2007Q4, against 40.8% in the euro area, and then fell 42.9% from the peak to their trough in 2013, which left them 27.5% above where they had started. Figure 1 plots the Spanish series alongside the other facts.

Three features of the cycle stand out. The first is consumption. Real private consumption rose 42.6% over the boom and fell 13.3% over the bust, a decline that is large by the standards of post-war European recessions and that tracks the housing cycle closely in timing. The second is homeownership. The homeownership rate was 79.8% before the boom, rose only to 80.5% at the peak, and had fallen to 77.7% by the trough. Spain entered the boom as a homeowner economy and left it as one. The adjustment over the cycle took place not in the share of households that owned, but in the value of what they owned

FIGURE 1. The Spanish boom and bust, 1996–2016.



NOTES: Top row: real house prices in Spain and the euro area, the real mortgage interest rate, and the median loan-to-value ratio at origination, which is annual and starts in 2004. Bottom row: real consumption, the employment rate, and real labour productivity per hour worked. House prices, consumption, employment and productivity are normalised to their 1996–1997 average. Vertical lines mark 1998Q1 and 2008Q1. Sources: Bank for International Settlements, Eurostat, Banco de España, and Bover, Torrado, and Villanueva (2019) for the median loan-to-value ratio.

and in the debt secured against it. Household mortgage debt rose from 0.23 to 0.62 of annual GDP, and the foreclosure rate peaked at 1.6% of mortgages. The boom, its financing and the bust are documented by Fernández-Villaverde, Garicano, and Santos (2013), Cano-Fuentes et al. (2013), Akin et al. (2014), Jimeno and Santos (2014), Estrada García and Saurina Salas (2016) and Martín, Moral-Benito, and Schmitz (2021), and the post-crisis adjustment of the housing market by Banco de España (2020) and Ghirelli, Leiva-León, and Urtasun (2022). The stability of the aggregate ownership rate hides a sharp fall among young households (Ferreira, Gálvez, and Pidkuyko 2024).

The third feature is construction, and it is as much a part of the Spanish story as credit. Housing starts rose from 323 thousand in 1997 to 760 thousand in 2006, and the volume of residential investment more than doubled; by the end of 2008 there were 614 thousand newly built dwellings without a buyer, and the volume of residential investment fell 52% from its 2007 peak to its 2013 trough. Residential investment went from 6.2% of GDP in 1997 to 11.7% in 2006. The censuses put the stock of dwellings at 20.9 million in 2001 and 25.2 million in 2011, but the number of households grew faster, from 14.2 to 18.1 million, so that dwellings per household fell from 1.47 to 1.39: much of what was built housed the households that immigration created, and much of the rest stood empty.

2.2. Credit Conditions and Labour Income

Three developments shaped the cycle: the fall in interest rates, the loosening and later tightening of lending standards, and the rise and fall of employment. Figure 1 plots them alongside house prices and consumption.

Interest rates. The convergence of Spanish rates on German ones in the run-up to monetary union cut the nominal mortgage rate from 5.50% at the end of 1997 to 2.25% at its lowest, and the real rate from an average of 4.06% over 1996–1997 to an average of 0.61% over the boom; the real risk-free rate fell from 3.27% to 0.26%. This is the single largest change in the household’s borrowing environment over the period, and it is exogenous to Spanish housing from the point of view of a small member of a currency union. Inflation, at 2.0% at the end of 1997 and an average of 3.0% over the boom, rose by about a point, so most of the fall in nominal rates was a fall in real rates. Rates then rose again into 2007, to 4.50% in nominal terms, which under adjustable-rate contracts raised the payment on every outstanding loan at the moment prices turned, and over the bust the real mortgage rate averaged 0.71% and the real risk-free rate 0.15%.

Labour income. The employment rate of the population aged 20 to 64 rose 27% between 1997 and its peak, an expansion without parallel in the euro area, and then fell back, ending 7% above its pre-boom level at the trough. Real labour productivity per hour was nearly flat over the boom, up 3%, and rose 10% over the bust as the least productive matches, concentrated in construction, were destroyed. Aggregate labour income per adult, the product of the two, rose by a factor of 1.31 to the peak and stood at 1.20 at its trough, where employment was 7% and productivity 13% above their 1997 levels.

Lending standards. Bover, Torrado, and Villanueva (2019) report the distribution of loan-to-value ratios on new Spanish mortgages, and Galán and Lamas (2019) the joint behaviour of loan-to-value, loan-to-income and debt-service ratios. Through the boom the median loan-to-value ratio stands at 100%: the representative buyer made no down payment. It falls abruptly to 80% after 2008, as lenders withdrew from high-leverage lending. Nothing in Spanish regulation required either number; both were lenders’ choices. There was no borrower-based macroprudential regulation in Spain over this period.

3. The Model

3.1. Environment

Time is discrete and a period is a quarter. The economy is populated by a unit continuum of infinitely lived households who differ in labour productivity, liquid wealth, housing, and mortgage debt. A

representative firm produces the non-durable consumption good, which is the numeraire. A competitive bank issues long-term, adjustable-rate, fully recourse mortgages, subject to the borrower-based caps that are the policy parameters of interest. Households save in a risk-free bond that trades at an exogenous rate on international markets; the economy is therefore small and open, and aggregate consumption need not equal aggregate output period by period. Owner-occupied houses are traded in a frictional market with directed search, intermediated by brokers. A construction sector converts the consumption good into housing quality and sells it to brokers; the housing stock depreciates and is maintained by whoever holds it. A rental technology supplies housing services elastically at a price proportional to the price of owner-occupied housing.

3.2. Households

Preferences. Households maximise

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{u(c_t, c_{h,t})^{1-\sigma}}{1-\sigma},$$

where $\beta \in (0, 1)$ is the discount factor, σ is the coefficient of relative risk aversion, c_t is consumption of the non-durable good and $c_{h,t}$ is consumption of housing services. Intratemporal utility is CES in the two goods,

$$u(c_t, c_{h,t}) = \left[\omega c_t^{\frac{\nu-1}{\nu}} + (1-\omega) c_{h,t}^{\frac{\nu-1}{\nu}} \right]^{\frac{\nu}{\nu-1}}, \quad (1)$$

with ν the intratemporal elasticity of substitution and ω the weight on the non-durable good.

Housing services. Households obtain housing services either by owning or by renting. A homeowner holds a house of quality $h_t \in \mathcal{H} = \{h_1, h_2, h_3\}$ and consumes $c_{h,t} = h_t$ every period. Housing depreciates at the rate ν per quarter, and whoever holds a house at the start of a quarter keeps it at its quality by purchasing νh_t units of quality that quarter at the price index p_t^h defined below; maintenance is due before consumption, falls on the start-of-quarter holder even if the house is sold during the quarter, and is first owed by a buyer in the quarter after the purchase. The discrete set $\mathcal{H}$ represents three tiers of the owner-occupied market. Households may hold at most one house at a time and may not let it. Renters purchase housing services on a spot market, choosing $c_{h,t} \in [0, h_r]$ with $h_r < h_1$. The restriction $h_r < h_1$ segments the two markets: the rental sector cannot deliver the quality that even the smallest owner-occupied house delivers, which is what generates a demand for ownership in equilibrium.

Endowments. Each household supplies one unit of labour inelastically, with idiosyncratic productivity $z_t e_t$. The persistent component z_t follows a finite-state Markov chain with transition matrix $\pi_z(z_{t+1} | z_t)$; the transitory component $e_t \in E \subset \mathbb{R}_+$ is drawn independently each period from $F(e)$. Labour income is $w_t = A_t z_t e_t$, where A_t is aggregate productivity.

3.3. The Consumption Good Producer

A representative firm operates a linear technology in labour under perfect competition. Let Φ_t denote the distribution of households over their individual states, defined in Section 3.10. Output is

$$Y_t = A_t \int z e d\Phi_t,$$

aggregate productivity times the aggregate supply of efficiency units.

3.4. The Housing Market

Houses are traded in a decentralised market with directed search, following Garriga and Hedlund (2020); the stock of owner-occupied housing evolves with construction and depreciation as set out in Section 3.5.

Structure. Buyers and sellers do not meet each other directly. A buyer transacts through a *buyer-broker* and a seller through a *seller-broker*, and brokers trade with one another frictionlessly at a per-unit-of-quality price p_t^h . Brokers may not carry inventories across periods, so they take no speculative position; they are passive market makers whose only role is to intermediate the flow of housing from sellers to buyers. Inventories nonetheless arise in equilibrium, because listings fail. This is the sense in which housing is illiquid in the model: a household that wishes to sell may not be able to.

The reason for the broker construct is tractability. A market in which heterogeneous buyers meet heterogeneous sellers requires each side to forecast the distribution of the other, and the joint distribution of income, liquid assets, housing and mortgage debt then becomes a state variable in every household's problem. Splitting the market in two—heterogeneous buyers against homogeneous brokers, homogeneous brokers against heterogeneous sellers—leaves only one heterogeneous side in each sub-problem, and delivers the block-recursive structure of Menzio and Shi (2010).

Directed search. Buyers choose a submarket indexed by a bid price and a house quality, $(p^{\text{bid}}, h) \in \mathbb{R}_+ \times \mathcal{H}$. Submarket tightness $\theta_b(p^{\text{bid}}, h)$ is the ratio of buyer-brokers to buyers there. A buyer meets a broker, and so purchases, with probability $\eta_b(\theta_b)$; a broker meets a buyer with probability $\alpha_b(\theta_b) = \eta_b(\theta_b)/\theta_b$. The function η_b is continuous and strictly increasing with $\eta_b(0) = 0$, while α_b is strictly decreasing and may exceed one, in which case a broker sells to more than one buyer. Matching is Cobb–Douglas with elasticity γ_b , so that

$$\eta_b(\theta_b) = \theta_b^{1-\gamma_b}, \quad \alpha_b(\theta_b) = \theta_b^{-\gamma_b}. \quad (2)$$

Sellers choose a submarket (p^{ask}, h) with tightness $\theta_s(p^{\text{ask}}, h)$, meeting probabilities $\eta_s(\theta_s) = \min\{\theta_s^{1-\gamma_s}, 1\}$ and $\alpha_s(\theta_s) = \eta_s(\theta_s)/\theta_s$, and elasticity γ_s ; the cap at one means that a seller who concedes enough sells with certainty, and the smallest such concession, a discount of $\kappa_s h$, defines a sure-sale

price that will matter below; no seller concedes more, since a larger discount buys no further probability. A buyer who fails to match remains a renter for the period. A seller who fails to match keeps the house and pays a utility cost ξ , which stands in for the cost of maintaining a listing.

Brokers and free entry. A buyer-broker who enters submarket (p^{bid}, h) pays $\kappa_b h$ and receives $p^{\text{bid}} - p_t^h h$ if matched; a seller-broker pays $\kappa_s h$ and receives $p_t^h h - p^{\text{ask}}$ if matched. Entry into every submarket is free, so

$$\kappa_b h \geq \alpha_b \left(\theta_b(p^{\text{bid}}, h) \right) (p^{\text{bid}} - p_t^h h), \quad \kappa_s h \geq \alpha_s \left(\theta_s(p^{\text{ask}}, h) \right) (p_t^h h - p^{\text{ask}}), \quad (3)$$

with equality in any submarket that is active. Combining (2) and (3) gives the tightness schedules

$$\theta_b(p^{\text{bid}}, h) = \left(\frac{p^{\text{bid}} - p_t^h h}{\kappa_b h} \right)^{\frac{1}{\gamma_b}}, \quad \theta_s(p^{\text{ask}}, h) = \left(\frac{p_t^h h - p^{\text{ask}}}{\kappa_s h} \right)^{\frac{1}{\gamma_s}}. \quad (4)$$

Block recursivity. Equation (4) is the key to the model's tractability and to its economics. The entire menu of trading probabilities available to households—how quickly one can buy at a given premium, how quickly one can sell at a given discount—depends on the cross-sectional distribution of households only through the single price index p_t^h . Households therefore need to forecast one number rather than a distribution, and p_t^h is pinned down by the requirement that the quality brokers buy from sellers, from the bank's reposessions and from builders equals the quality they sell to buyers and to owners and the bank for maintenance. Builders and maintaining holders trade with brokers at p_t^h and nowhere else, which is why adding a construction sector leaves the structure intact.

The economics is that liquidity is endogenous and procyclical without being a free parameter. A buyer must post $p^{\text{bid}} > p_t^h h$ to be served at all, and the premium buys a higher trading probability; a seller must accept $p^{\text{ask}} < p_t^h h$, and the discount buys speed. When the economy expands, more households wish to buy and fewer are forced to sell: at the going p_t^h , buyer submarkets become tight and seller submarkets slack, sellers trade quickly at small discounts, and p_t^h rises. When it contracts, the same mechanism runs in reverse and the expected time on the market lengthens. A household that receives a bad income shock in a slump cannot count on converting its house into cash, and must cut consumption instead.

3.5. Housing Supply

A competitive construction sector converts the consumption good into housing quality. Building I_t units of quality in a quarter costs $C(I_t)$ units of the good, with C increasing and convex, and builders sell what they build to brokers at the frictionless price index p_t^h . Builders therefore produce up to the point where marginal cost equals the price, and with an isoelastic cost function the supply of new

housing is

$$I_t = \nu H_{ss} \left(\frac{p_t^h}{p_{ss}^h} \right)^\lambda, \quad C(I_t) = \frac{\lambda}{1 + \lambda} p_t^h I_t, \quad (5)$$

where λ is the price elasticity of new construction and the subscript ss denotes the pre-boom stationary equilibrium. The aggregate stock of housing quality H_t , held by owners and by the bank, evolves as

$$H_{t+1} = (1 - \nu) H_t + I_t, \quad (6)$$

so that in a stationary equilibrium construction replaces depreciation, $I = \nu H$, and over a boom the stock grows by the excess of building over depreciation and carries that excess into the bust. The rents that builders earn above cost, $p_t^h I_t - C(I_t)$, accrue to a fixed factor owned abroad, so that household budgets are unaffected; the cost of building is a use of the economy's output, and the resource constraint is stated in Section 3.10.

Two things the sector does not do should be said at the outset. It does not create jobs: labour income is taken from the data, so the employment that Spain's construction boom created and its bust destroyed is in the model as an input, not as a consequence of building, and a cap that reduces construction in the model leaves labour income unchanged. On the income side, then, the counterfactuals understate what a cap would have done to the bust and understate what it would have cost in the boom. And the sector holds no inventories: what builders build is absorbed by brokers within the quarter, so the unsold new dwellings that accumulated in Spain after 2007 have no counterpart; the overhang the model carries is the stock that owners list and fail to sell.

3.6. The Rental Technology

Households can exchange p_t^r units of the numeraire for one unit of housing services, up to the limit h_r . The rental price is proportional to the price index in the owner-occupied market,

$$p_t^r = \zeta p_t^h,$$

where ζ is the rent-to-price ratio. This is an abstraction: it dispenses with landlords, with the conversion of owner-occupied units into rental units, and with rental market dynamics generally, and makes the rental sector a perfectly elastic outside option. The abstraction is defensible for Spain over this period, where the rental sector was small and its price moved closely with the owner-occupied index, and it is conservative for the question at hand: a rental market that expanded during the bust would have provided households with an additional margin of adjustment and would have made the consumption collapse smaller.

3.7. Saving and Mortgages

Saving. Households save in a one-period risk-free bond $b_{t+1} \geq 0$ that trades at the exogenous rate r_t^b set on international markets. Because the rate is exogenous and borrowing in the bond is not allowed, the bond is the households' liquid buffer and the channel through which the household sector as a whole borrows from, or lends to, the rest of the world.

The mortgage contract. A competitive bank offers collateralised borrowing through long-term mortgages. The contract mirrors the Spanish one over the sample period and differs in three respects from the U.S. contract usually modelled. It carries an adjustable rate, so the payment on an outstanding balance moves with the market rate r_t^m rather than being fixed at origination; over 90 per cent of Spanish mortgages outstanding in 2007 were variable rate. It has no fixed maturity in the model: the borrower chooses the pace of amortisation subject to a minimum payment. And it is a full-recourse loan: a borrower whose house is repossessed and sold for less than the outstanding balance remains liable for the shortfall. A mortgage is therefore fully described by its outstanding balance m_t .

Origination. A household with cash on hand y_t , persistent productivity z_t , and a house of quality h_t that takes out a loan of face value m_{t+1} receives $q_t(y_t, z_t, h_t, m_{t+1}) m_{t+1} < m_{t+1}$ units of the numeraire. The price $q_t(\cdot)$ prices the borrower's risk and is set so that the bank makes zero expected profit loan by loan; it is derived in Section 3.9. The borrower pays a proportional origination cost $\mu q_t m_{t+1}$, and the loan is subject to three borrower-based constraints at origination:

$$q_t(y_t, z_t, h_t, m_{t+1}) m_{t+1} \leq \vartheta p_t^h h_t, \quad (7)$$

$$q_t(y_t, z_t, h_t, m_{t+1}) m_{t+1} \leq \delta \cdot 4A_t z_t e_t, \quad (8)$$

$$a(i_t^m) m_{t+1} \leq \psi A_t z_t e_t. \quad (9)$$

The first is a loan-to-value cap, with ϑ the maximum ratio of the loan to the value of the house. The second is a debt-to-income cap, with δ the maximum ratio of the loan to annual earnings. The third is a payment-to-income cap: $a(i_t^m)$ is the quarterly payment per unit of principal on a standard 25-year amortising loan at the nominal mortgage rate i_t^m , and ψ is the maximum ratio of that payment to quarterly earnings. The caps apply to earnings, not to cash on hand, as they do in practice. The nominal rate is the one households face in the data; the model is written in real terms, with $1 + r_t^m = (1 + i_t^m)/(1 + \text{inflation}_t)$ at quarterly frequency, the annual rates in the data having been converted once to quarterly rates before any of this is applied, so the payment cap is a rule indexed to an observed nominal rate rather than a feature of the real contract, which is what a supervisor's rule is. In the benchmark economy only (7) is in place, at the lax level of Section 4; the policy experiments of Section 6 set ϑ , δ and ψ one at a time and in combination.

Repayment and refinancing. The borrower chooses a payment l_t each period. The contract requires only that the balance not grow,

$$m_{t+1} = (m_t - l_t)(1 + r_t^m) \leq m_t \iff l_t \geq \frac{r_t^m}{1 + r_t^m} m_t, \quad (10)$$

so the minimum payment is the interest due, and it moves with the market rate. A borrower may instead extinguish the balance and originate a new loan, paying the origination cost again; this is how equity is extracted, since the new balance may exceed the old.

Delinquency, repossession and deficiency. A homeowner who does not make the minimum payment, does not refinance, and has not sold the house is delinquent. With probability φ the bank repossesses. It then sells the house, incurring a proportional loss χ on the proceeds, and the household becomes a renter owing the shortfall

$$d_t = \max\{0, m_t - (1 - \chi) p_t^h h_t\}. \quad (11)$$

Under the deficiency judgment a fraction τ of the household's income is withheld each period until d_t is extinguished; the debt cannot be discharged. With probability $1 - \varphi$ the bank does not repossess: the missed payment is ignored and the balance m_t carries into the next period.

This is the contractual feature that matters most for the Spanish application. Under U.S.-style non-recourse lending, default caps the borrower's loss at the value of the house and the option to walk away is a form of insurance. Under Spanish full recourse it is not: a household that loses its house in the bust enters a long period during which 30% of its income is withheld, and its consumption falls accordingly. The households who supply the sharp part of the consumption collapse in Section 5 are precisely these, and the welfare gain from macroprudential regulation in Section 6.4 is largely the gain from there being fewer of them.


3.8. The Household's Problem

The state of a renter is (y_t, z_t) , where $y_t = A_t z_t e_t + b_t$ is cash on hand; the state of a homeowner is (y_t, z_t, m_t, h_t) ; the state of a household under a deficiency judgment is (y_t, z_t, d_t) . Within a period, decisions are taken in the order set out in Figure 2: productivity is drawn, owners decide whether to list, owners decide whether to pay, refinance or go delinquent, non-owners decide whether to bid, and finally consumption and saving are chosen.

Tenure choice. A household that begins the period without a house chooses between remaining a renter and entering the housing market as a bidder:

$$W_t^{\text{no}}(y_t, z_t) = \max\{V_t^r(y_t, z_t), W_t^{\text{bid}}(y_t, z_t)\}.$$

FIGURE 2. Timing within a period.



NOTES: Dashed lines mark events; arrows mark household decisions. Listing and bidding are choices of a submarket, that is, of a price and—for buyers—a house quality.

The renter. A renter chooses non-durable consumption, housing services and saving:

$$V_t^r(y_t, z_t) = \max_{\substack{c_t \geq 0, b_{t+1} \geq 0 \\ c_{h,t} \in [0, h_r]}} \left\{ \frac{u(c_t, c_{h,t})^{1-\sigma}}{1-\sigma} + \beta \mathbb{E}_{z_{t+1}|z_t} [W_{t+1}^{\text{no}}(y_{t+1}, z_{t+1})] \right\}$$

subject to $c_t + p_t^r c_{h,t} + b_{t+1}/(1+r_t^b) \leq y_t$ and $y_{t+1} = A_{t+1}z_{t+1}e_{t+1} + b_{t+1}$.

The bidder. A household that enters the market chooses a bid price and a quality, trading a higher price against a higher probability of being served:

$$W_t^{\text{bid}}(y_t, z_t) = \max_{p_t^{\text{bid}} \geq 0, h_t \in \mathcal{H}} \left\{ \eta_b(\theta_b(p_t^{\text{bid}}, h_t)) V_t^b(y_t, z_t, h_t) + [1 - \eta_b(\theta_b(p_t^{\text{bid}}, h_t))] V_t^r(y_t, z_t) \right\}.$$

The buyer. A household that is matched buys the house, takes out a mortgage, and consumes:

$$V_t^b(y_t, z_t, h_t) = \max_{c_t \geq 0, b_{t+1} \geq 0, m_{t+1} \geq 0} \left\{ \frac{u(c_t, h_t)^{1-\sigma}}{1-\sigma} + \beta \mathbb{E}_{z_{t+1}|z_t} [W_{t+1}^o(y_{t+1}, z_{t+1}, m_{t+1}, h_t)] \right\}$$

subject to the budget constraint, the origination constraints (7)–(9), and the law of motion of cash on hand,

$$c_t + p_t^{\text{bid}} + \frac{b_{t+1}}{1+r_t^b} \leq y_t + (1-\mu) q_t(y_t, z_t, h_t, m_{t+1}) m_{t+1}, \quad y_{t+1} = A_{t+1}z_{t+1}e_{t+1} + b_{t+1}.$$

Listing choice. A household that begins the period owning chooses whether to list:

$$W_t^o(y_t, z_t, m_t, h_t) = \max \left\{ W_t^{\text{ask}}(y_t, z_t, m_t, h_t), W_t^h(y_t, z_t, m_t, h_t) \right\},$$

where the value of listing at price p_t^{ask} is

$$W_t^{\text{ask}}(\cdot) = \max_{p_t^{\text{ask}} \geq m_t - \tilde{y}_t} \left\{ \eta_s(\theta_s(p_t^{\text{ask}}, h_t)) \left[V_t^r(\tilde{y}_t + p_t^{\text{ask}} - m_t, z_t) + W_t^{\text{bid}}(\tilde{y}_t + p_t^{\text{ask}} - m_t, z_t) \right] \right\}$$

$$+ \left[1 - \eta_s \left(\theta_s(p_t^{\text{ask}}, h_t) \right) \right] \left[W_t^h(y_t, z_t, m_t, h_t) - \xi \right] \}.$$

Here $\tilde{y}_t = y_t - v p_t^h h_t$ is cash on hand net of the quarter's maintenance, which the seller owes as the start-of-quarter holder. The restriction $p_t^{\text{ask}} \geq m_t - \tilde{y}_t$ is the requirement that the sale proceeds together with cash on hand cover the outstanding balance: a household deep enough underwater cannot post a price that clears its mortgage, and is therefore locked into its house. This is the collateral channel in its sharpest form, and it is why the distribution of leverage entering the bust—the object macroprudential policy acts on—determines how much consumption falls.

It is also where the inefficiency that justifies a cap resides, and it is worth stating in plain terms. When a household takes a larger mortgage, it adds, in a small way, to the number of households that will be forced to sell or to default when prices fall, and so to the fall in p_t^h itself. That fall tightens the constraint $p_t^{\text{ask}} \geq m_t - \tilde{y}_t$ for every *other* indebted household: houses that would have covered their debts no longer do, their owners cannot sell, and they default. The borrower does not pay for this harm, and the bank, which prices the borrower's own default in q_t , does not charge for it either, because it prices against the equilibrium path of p_t^h that no single loan moves. In the states in which the listing constraint binds, then, households borrow more than is good for all of them together, and a cap makes all of them borrow less at once, which none would do alone; whether the gain outweighs what the cap costs the households it excludes, and the lower prices it hands to sellers in the boom, is the question the welfare results of Section 6 answer. This is the pecuniary externality of Lorenzoni (2008), Bianchi (2011) and Dávila and Korinek (2018), operating through a price that sits inside a constraint rather than through a collateral limit on borrowing. The search market is not the source of it: given the states households bring to the market, the allocation across submarkets under free entry is efficient in the sense of Moen (1997). What the friction does is make the externality bite harder, by turning a fall in the price into an inability to sell at all.

The homeowner. A homeowner who does not sell chooses between making a payment, refinancing, and going delinquent:

$$W_t^h(y_t, z_t, m_t, h_t) = \max \left\{ V_t^p(\cdot), V_t^m(\cdot), \varphi V_t^d(y_t, z_t, d_t) + (1 - \varphi) V_t^s(\cdot) \right\},$$

where d_t is given by (11). The value of making a payment is

$$V_t^p(\cdot) = \max_{c_t \geq 0, b_{t+1} \geq 0, l_t} \left\{ \frac{u(c_t, h_t)^{1-\sigma}}{1-\sigma} + \beta \mathbb{E}_{z_{t+1}|z_t} \left[W_{t+1}^o(y_{t+1}, z_{t+1}, m_{t+1}, h_t) \right] \right\}$$

subject to $c_t + v p_t^h h_t + b_{t+1}/(1+r_t^b) + l_t \leq y_t$, the law of motion $m_{t+1} = (m_t - l_t)(1+r_t^m)$, and the minimum payment (10). The value of refinancing, V_t^m , replaces these by

$$c_t + v p_t^h h_t + \frac{b_{t+1}}{1+r_t^b} + m_t \leq y_t + (1-\mu) q_t(y_t, z_t, h_t, m_{t+1}) m_{t+1}$$

and the origination constraints (7)–(9). A homeowner who skips a payment and is not repossessed consumes out of cash on hand net of maintenance and carries $m_{t+1} = m_t$ forward; if cash on hand does not cover maintenance, the bank advances the shortfall and adds it to the balance, as it would an escrowed charge, so that delinquency remains feasible in every subsequent quarter and the advance is priced into q_t through the larger balance. A household that is repossessed becomes a renter subject to the garnishment, solving

$$V_t^d(y_t, z_t, d_t) = \max_{\substack{c_t \geq 0, b_{t+1} \geq 0 \\ c_{h,t} \in [0, h_r]}} \left\{ \frac{u(c_t, c_{h,t})^{1-\sigma}}{1-\sigma} + \beta \mathbb{E}_{z_{t+1}|z_t} \left[\mathbf{1}\{d_{t+1} > 0\} V_{t+1}^d(y_{t+1}, z_{t+1}, d_{t+1}) \right. \right. \\ \left. \left. + \mathbf{1}\{d_{t+1} = 0\} W_{t+1}^{\text{no}}(y_{t+1}, z_{t+1}) \right] \right\}$$

subject to $c_t + p_t^r c_{h,t} + b_{t+1}/(1 + r_t^b) \leq (1 - \tau) y_t$ and $d_{t+1} = \max\{0, (d_t - \tau y_t)(1 + r_t^m)\}$.

Owners who cannot pay maintenance. The test is made at the listing stage, with the cash on hand known then, against the better of the payment and refinancing branches. An owner for whom that branch leaves no non-negative consumption after maintenance sells at the sure-sale price $p_t^h h_t - \kappa_s h_t$, which exists because η_s is capped at one; this is always feasible for an owner without a mortgage, and for an indebted owner whenever $p_t^h h_t - \kappa_s h_t \geq m_t - \tilde{y}_t$. An indebted owner for whom even the sure sale would not clear the balance goes delinquent, and the unpaid maintenance is advanced by the bank as described above. A house the bank holds at the start of a quarter is maintained by the bank. Every branch of the household's problem is therefore well defined at every state, and total maintenance in a quarter is v times the start-of-quarter stock.

3.9. Mortgage Pricing

The bank is competitive and funds itself at the risk-free rate plus a servicing spread, so the mortgage rate satisfies $1 + r_t^m = (1 + r_t^b)(1 + \phi_t)$. The model takes the nominal mortgage rate, the nominal risk-free rate and inflation from the data and derives the real rates from them, so the spread ϕ_t is the residual between the two real rates and moves over the transition as it did in the data, remaining positive in every quarter; in the stationary equilibria it is constant. It prices each loan so that its expected discounted repayment equals the amount disbursed. Let starred variables denote the borrower's policies at $t + 1$ evaluated at the state implied by (y_t, z_t, h_t, m_{t+1}) , and let η_{t+1}^{s*} be the probability that the borrower's chosen listing succeeds. Then

$$q_t m_{t+1} = \frac{1}{1 + r_{t+1}^m} \mathbb{E}_t \left[\underbrace{\eta_{t+1}^{s*} m_{t+1}}_{\text{sale}} + (1 - \eta_{t+1}^{s*}) \left\{ \underbrace{\mathbf{1}_{t+1}^{m*} m_{t+1}}_{\text{refinance}} + \underbrace{\mathbf{1}_{t+1}^{p*} (l_{t+1}^* + q_{t+1}^c m_{t+2}^*)}_{\text{payment}} \right\} \right]$$

$$\begin{aligned}
& + \mathbf{1}_{t+1}^{d*} \left[\underbrace{\varphi \left((1 - \chi) p_{t+1}^h h_t - M_{t+1} + Q_{t+1}^{dj} \right)}_{\text{repossession}} \right. \\
& \quad \left. + (1 - \varphi) \underbrace{\left(q_{t+1}^s (m_{t+1} + a_{t+1}) - a_{t+1} \right)}_{\text{skipped payment}} \right] \Bigg\}, \tag{12}
\end{aligned}$$

where the loan is repaid in full when the house is sold; $\mathbf{1}^{m*}$, $\mathbf{1}^{p*}$ and $\mathbf{1}^{d*}$ are the indicators of the refinancing, payment and delinquency branches of W_{t+1}^h , and in the payment branch the bank receives the payment plus the continuation value of the loan; q_{t+1}^c marks to market the continuation of a loan with updated balance m_{t+2}^* ; $a_{t+1} = \max\{0, \nu p_{t+1}^h h_t - y_{t+1}\}$ is the maintenance the bank advances to a delinquent borrower who cannot pay it, which it adds to the balance, and q_{t+1}^s marks that balance to market; M_{t+1} is the maintenance the bank expects to pay while it holds a repossessed house; and Q_{t+1}^{dj} is the expected discounted value of the garnishment stream on the deficiency d_{t+1} , which is bounded above by d_{t+1} itself. The bank's maintenance is separate from the resale loss χ , which is a real cost of selling through the market. Appendix B states the recursion for each of these objects.

Two properties of (12) matter for what follows. First, the price falls with expected illiquidity: a lower η_{t+1}^{s*} removes probability from the branch in which the loan is repaid in full and adds it to the branches in which it is not, and it also lengthens the time the bank expects to hold a repossessed house. Credit conditions therefore tighten endogenously when the housing market becomes illiquid, without any change in lender behaviour. Second, full recourse raises the price relative to a non-recourse contract, because $Q_{t+1}^{dj} > 0$: Spanish borrowers could borrow more cheaply for a given loan size than U.S. borrowers could, and they paid for it after 2008.

3.10. Equilibrium

Stationary equilibrium. Given constant $(A, r^b, r^m, i^m, \vartheta, \delta, \psi)$ and the supply parameters (λ, ν) with the pre-boom (H_{ss}, p_{ss}^h) that normalise (5), a stationary recursive equilibrium consists of value functions $\{W^{no}, V^r, W^{bid}, V^b, W^o, W^{ask}, W^h, V^p, V^m, V^s, V^d\}$ and associated policies; a mortgage pricing function q ; tightness schedules (θ_b, θ_s) ; prices (p^h, p^r) ; a housing stock H and construction I ; a distribution Φ over household states; and the bank's holdings of repossessed houses by quality, which grow with repossessions and shrink with the bank's sales and are part of H , such that households optimise; q satisfies (12); (θ_b, θ_s) satisfy (4); $p^r = \zeta p^h$; I satisfies (5) and replaces depreciation, $I = \nu H$; Φ is invariant under the policies; and p^h clears the housing market, in the sense that the quality brokers purchase from sellers, from the bank's repossessions and from builders equals the quality they sell to buyers and to owners and the bank for maintenance.

Resource constraint. Because the economy is open, its budget is a law of motion for its net foreign asset position, the households' bonds net of the bank's foreign funding of the mortgage stock:

$$\frac{NFA_{t+1}}{1 + r_t^b} = NFA_t + Y_t - C_t - X_t - \Pi_t^b,$$

where C_t is aggregate non-durable consumption, $\Pi_t^b = p_t^h I_t - C(I_t)$ is the builders' rent paid abroad, and X_t collects the real resources the housing and credit system absorbs: the cost of building $C(I_t)$, the resources that supply rental services, the brokers' entry costs, the origination cost μ , the bank's loss χ on repossessed houses, and the servicing spread. Maintenance, garnishment and deficiency payments are transfers between households, builders and the bank and net out. The bank breaks even in expectation; the gains and losses it makes when a shock arrives unanticipated accrue to the foreign funders of the mortgage stock and appear in the net foreign asset position. GDP is Y_t ; model consumption is C_t , which excludes construction and housing services and so corresponds to non-housing private consumption in the national accounts.

Transition equilibrium. The experiments in Sections 5 and 6 are perfect-foresight transitions between stationary equilibria, punctuated by unanticipated shocks. Given an initial distribution Φ_0 and stock H_0 , a terminal stationary equilibrium, and sequences $\{A_t, r_t^b, r_t^m, i_t^m, \vartheta_t, \delta_t, \psi_t\}_{t \geq 0}$, a transition equilibrium is a sequence of value functions, policies, pricing functions, tightness schedules, prices, construction, stocks and distributions satisfying the stationary conditions period by period, with Φ_{t+1} generated from Φ_t by the date- t policies, H_{t+1} by (6), and p_t^h clearing the housing market at each date; the terminal stock is the one the path leads to. At an unanticipated shock, households solve a new problem taking the prevailing distribution and stock as given, having assigned zero probability to the event. In the terminal stationary equilibrium inflation is set to two per cent a year, which fixes the nominal rate that the payment cap is assessed at once the data end. Appendix C describes the computation.

4. Calibration

The model is calibrated to the Spanish economy in 1997Q4, immediately before the boom. Wherever possible, target moments are measured at that date; where the series does not reach back that far, the nearest available observation is used. Parameters are assigned in two groups: those set outside the model, and those chosen jointly so that the stationary equilibrium reproduces a set of target moments. The unit of account is mean annual earnings, which the aggregate productivity parameter A normalises to one.

4.1. External Parameters

Table 1 lists the externally assigned parameters.

Earnings. The persistent and transitory components of labour productivity are taken from Guner, López-Segovia, and Ramos (2020), who estimate the process on Spanish survey earnings, with tax records for the top of the distribution. Their estimates are annual and the model period is a quarter. The persistent component is converted by taking the fourth root of its annual autocorrelation, 0.983, which gives 0.9957 a quarter, and by scaling the standard deviation of its innovations from 0.072 to 0.036 so that the unconditional dispersion of z is unchanged. The transitory component is redrawn each quarter with its annual standard deviation, 0.299; summed over a year it then accounts for a quarter of the annual transitory variance, a choice that errs on the side of less transitory risk. The persistent component is discretised into a three-state Markov chain by the method of Rouwenhorst (1995), and a fourth state is appended to capture the top one per cent of earners, whose behaviour matters for the upper tail of the wealth distribution and hence for the housing market. The probability of remaining in the top state, $\pi_{4,4} = 0.858$, is taken from tax records; the probability of entering it from the third state is set so that one per cent of households occupy it, giving $\pi_{3,4} = 0.006$; direct transitions into the top state from the two lowest states are ruled out. Appendix D reports an alternative process estimated directly on the Spanish EU-SILC and shows that the results are unaffected.

Preferences. There is no direct estimate of the intratemporal elasticity of substitution between housing and non-durable consumption for Spain, so I use the estimate of Flavin and Nakagawa (2008) for the United States, $\nu = 0.13$. Risk aversion is set to the standard $\sigma = 2$.

Rental market. The rent-to-price ratio is $\zeta = 0.022$, computed from the 2002 wave of the Spanish Survey of Household Finances (*Encuesta Financiera de las Familias*) as the ratio of reported rent to reported dwelling value.

Financial markets. The model takes the nominal one-year Treasury rate, the nominal mortgage reference rate and consumer price inflation from the data and deflates them exactly. The real risk-free rate in the pre-boom stationary equilibrium is the 1996–1997 average, 3.27%, and the real mortgage rate the corresponding average, 4.06%; the servicing spread ϕ is the residual between the two, and over the transition it is the residual between the two observed series.

Housing supply. Housing depreciates at 2.0% a year, $\nu = 0.50\%$ per quarter. This is above the measured rate. The Fundación BBVA–Ivie capital stock series puts the geometric depreciation rate of dwellings at 1.14% a year in its current vintage, which combines an average service life of 80 years with a declining-balance factor of 0.91, and used 3.33%, a double-declining balance over 60 years, in its earlier ones (Mas Ivars and Pérez García 2021). The rate is set higher because construction in a stationary economy only replaces depreciation, whereas part of Spanish residential investment in 1997 added to a growing

TABLE 1. External parameters.

Parameter	Symbol	Value	Source
<i>Earnings</i>			
Persistence, quarterly (annual 0.983)	ρ_z	0.9957	Guner, López-Segovia, and Ramos (2020)
Innovation s.d., quarterly (annual 0.072)	σ_z	0.036	Guner, López-Segovia, and Ramos (2020)
Transitory s.d., redrawn each quarter	σ_e	0.299	Guner, López-Segovia, and Ramos (2020)
Persistence in the top 1%	$\pi_{4,4}$	0.858	Spanish tax records
Transition into the top 1%	$\pi_{3,4}$	0.006	Set so 1% of households have $z = z_4$
<i>Preferences and housing</i>			
Intratemporal elasticity of substitution	ν	0.13	Flavin and Nakagawa (2008)
Risk aversion	σ	2	Standard
Rent-to-price ratio	ζ	0.022	EFF 2002
<i>Financial markets and policy</i>			
Real risk-free rate (annual)	r^b	3.27%	Banco de España, 1996–1997 average
Real mortgage rate (annual)	r^m	4.06%	Banco de España, 1996–1997 average
Housing depreciation (annual)	υ	2.0%	Mas Ivars and Pérez García (2021); see text
Maximum loan-to-value ratio	ϑ	100%	Bover, Torrado, and Villanueva (2019)
Probability of repossession	φ	15.4%	CGPJ, foreclosure duration
Garnishment rate, deficiency judgment	τ	30%	Civil Procedure Act, Art. 607

NOTES: Rates are annualised; the model period is a quarter. The unit of account is mean annual earnings. The maximum loan-to-value ratio is the upper envelope of the origination distribution, not a regulatory ceiling; it does not bind for the median borrower in the benchmark economy.

stock: at the measured rate, with λ and the housing-wealth target, construction would absorb 2.6% of output before the boom against 6.2% in the data; at 2.0% it absorbs about three-quarters of the data share. The elasticity of new construction λ is set internally, as described below.

Policy and legal environment. The maximum loan-to-value ratio in the benchmark is $\vartheta = 100\%$. This is not a regulatory ceiling—there was none—but the upper envelope of the loan-to-value distribution reported by Bover, Torrado, and Villanueva (2019), and it is not binding for the median borrower: households in the benchmark economy are not required to make a down payment. There was no income-based cap either, so δ and ψ are unbounded in the benchmark. Spanish mortgages carry full recourse: a household that loses its house remains liable for the shortfall (see Galán, Lamas, and Vegas 2022, for an analysis of recourse in Spanish mortgage lending). The probability of repossession per quarter of delinquency, $\varphi = 15.4\%$, is set so that the expected time from the first missed payment to repossession under a geometric distribution matches the roughly 20 months that the General Council of the Judiciary reports for the judicial foreclosure procedure. The garnishment rate under a deficiency judgment, $\tau = 30\%$, is the rate that Article 607 of the Civil Procedure Act applies to the first tranche of earnings above the minimum wage; the statutory schedule rises by tranche to 90%, and the single rate is a simplification.

4.2. Internally Calibrated Parameters

Thirteen parameters are set jointly: twelve so that the stationary equilibrium matches twelve moments, and the construction elasticity so that the benchmark transition matches one. Table 2 lists them together with the moment each is most closely associated with, and the fit. Although all parameters are determined jointly, the mapping between them and the targets is close to one to one, which is what the pairing in the table records.

The elasticity of construction λ is the one parameter that the pre-boom economy cannot identify, because in a stationary equilibrium building only replaces depreciation. It is set so that new construction in the model rises over the boom by the same factor as the volume of residential investment in the national accounts, which rose by a factor of 2.16 between 1997 and its 2007 peak; since building responds to the price index with elasticity λ and the benchmark price boom is of almost the same size, this gives $\lambda = 1.01$. Two implications are then checks rather than targets. The resources that construction absorbs in the pre-boom equilibrium, $[\lambda/(1 + \lambda)] v p^h H/Y$, are pinned by λ , v and the housing-wealth target, and come to 4.6% of GDP against 6.2% in the data; and the rise in that share over the boom, which the model tracks at constant prices and overstates at current prices because its construction cost moves with the house price, there being no land, is reported in Section 5.2.

The three house qualities are anchored by h_1 , which is set so that the value of the smallest owner-occupied house matches the tenth percentile of dwelling values in the 2002 EFF; the two higher tiers are spaced geometrically with ratio 1.65, giving $\mathcal{H} = \{3.28, 5.41, 8.93\}$. The ceiling on rental quality, h_r , is the parameter that governs how attractive renting is relative to owning, and it is pinned down by the homeownership rate. The discount factor β governs how much households save before buying and therefore the median loan-to-value ratio at origination; the consumption weight ω governs the split of the portfolio between housing and other wealth and is identified by mean housing wealth; the ratio z_4/z_3 governs the upper tail and is identified by mean net worth.

The four search parameters are identified by the four dimensions of housing market activity. The matching elasticities γ_b and γ_s govern how steeply the trading probability rises with the price concession and so determine expected search and listing times, which are targeted to the durations reported by Fotocasa Research. The broker entry costs κ_b and κ_s determine how large a concession is ever worth making, and so are identified by the maximum bid premium and the maximum listing discount observed in matched listing and transaction data. No Spanish estimate of either moment exists, so the targets, 11.0% and 2.5%, are the values used in the directed-search housing literature, and Section 7.1 reports how much the results depend on them. The utility cost of a failed listing ξ governs how readily households list speculatively and is identified by the share of the stock on the market. Finally, the origination cost μ is identified by the flow of new mortgages, and the bank's resale discount χ by the foreclosure rate, since a larger discount makes repossession more costly to the bank and is passed back into the pricing function.

TABLE 2. Internally calibrated parameters and targeted moments.

Parameter	Symbol	Value	Target	Data	Model
Maximum rental quality	h_r	1.25	Homeownership rate	79.5%	79.7%
Minimum owned quality	h_l	3.28	Starter house value	1.24	1.39
Productivity at the top	z_4/z_3	5.50	Mean net worth	3.14	3.71
Consumption weight	ω	0.816	Mean housing wealth	4.37	4.56
Discount factor	β	0.974	Median LTV at origination	80.0%	76.8%
Matching elasticity, buyers	γ_b	0.906	Buyer search (weeks)	5.7–11.5	7.8
Matching elasticity, sellers	γ_s	0.345	Listing time (months)	2.9–6.1	4.8
Cost of a failed listing	ξ	0.0013	Share of stock for sale	2.7%	3.1%
Broker entry cost, buyers	κ_b	0.118	Maximum bid premium	11.0%	10.0%
Broker entry cost, sellers	κ_s	0.024	Maximum list discount	2.5%	4.2%
Mortgage origination cost	μ	0.019	New mortgages	2.95%	4.32%
Bank's resale discount	χ	0.137	Foreclosure rate	1.1%	1.1%
Elasticity of construction	λ	1.01	Rise in new housing over the boom	2.16	2.16

NOTES: All parameters are set jointly; the pairing records which moment each parameter is most closely associated with. Sources: INE for the homeownership rate, new mortgages and the volume of residential investment; the 2002 EFF for the starter house value, net worth, housing wealth and the median loan-to-value ratio; Fotocasa Research for buyer search and listing times; Idealista for the share of the stock for sale; the directed-search housing literature for the maximum bid premium and list discount, for which no Spanish estimate exists; CGPJ for the foreclosure rate. Values expressed in units of mean annual earnings are the starter house value, mean net worth and mean housing wealth. Buyer search and listing times are reported as ranges because the source reports them by region. The rise in new housing is the ratio of construction at its peak to construction in 1997Q4, matched to the volume index of gross fixed capital formation in dwellings; it is the one moment taken from the transition rather than from the stationary equilibrium.

4.3. Fit

The calibrated economy is within a tenth of its targets for the homeownership rate, housing wealth, the bid premium and the foreclosure rate, and inside the reported ranges for search and listing times; it overshoots the starter house value, mean net worth, the share of the stock for sale, the list discount and the flow of new mortgages by more. Two misses matter for what follows. The median loan-to-value ratio at origination is 76.8% against 80.0% in the data: the model does not quite reproduce how little Spanish households saved before buying, and this makes the benchmark economy slightly less levered entering the boom than the data, which biases the estimated effect of macroprudential policy downwards. The flow of new mortgages is 4.32% against 2.95%: the model generates somewhat too much refinancing at the margin.

Table 3 reports moments that were not targeted. The model reproduces the concentration of wealth reasonably well, with a net worth Gini of 0.53 against 0.57, and it reproduces the flat homeownership profile across the income distribution that distinguishes Spain from most other advanced economies: 61.2% of the bottom income quintile own their home in the model, against 63.8% in the data. It also reproduces the level of household indebtedness at the starting point, at 0.26 of annual GDP against 0.23, which matters because the counterfactuals act on the growth of that number.

TABLE 3. Untargeted moments, 1997–2002.

Moment	Data	Model
Gini coefficient, net worth	0.57	0.53
Gini coefficient, housing wealth	0.42	0.39
Homeownership rate, bottom income quintile	63.8%	61.2%
Homeownership rate, top income quintile	90.1%	92.4%
Share of households with mortgage debt	27.4%	25.8%
Mean loan-to-value ratio, indebted owners	41.6%	38.9%
Housing wealth over net worth	0.79	0.81
Sales over the owner-occupied stock (annual)	3.4%	3.0%
Mortgage delinquency rate	0.9%	1.0%
Household mortgage debt over annual GDP	0.23	0.26
Residential investment over GDP, 1997	6.2%	4.6%

NOTES: Distributional moments are computed from the 2002 wave of the EFF, the earliest available, and compared with the model's stationary distribution. Aggregate moments are 1997Q4 or the nearest available date. Sources: EFF, INE, Banco de España, CGPJ.

5. The Boom and Bust in the Model

5.1. The Experiment


The economy starts in the stationary equilibrium of Section 4, which represents 1997Q4. It is then hit by two unanticipated sequences of shocks, the first at 1998Q1 and the second at 2008Q1. Households have perfect foresight between the two dates and assign zero probability to the second.

The boom shock has two components. Aggregate productivity A_t rises gradually toward a new, higher stationary level, reproducing the combined effect of the 27% rise in employment and the nearly flat path of labour productivity per hour; at the 2007 peak A_t stands at 1.31 times its pre-boom value. Simultaneously the real mortgage rate falls from an average of 4.06% to an average of 0.61% and the real risk-free rate from 3.27% to 0.26%, on the paths observed in the data, with the nominal rates and inflation that produce them; the nominal mortgage rate matters in its own right only once a payment cap is in place. Households treat the new regime as permanent.

The bust shock also has two components. Aggregate productivity declines gradually toward a new, lower stationary level, at 1.20 times its pre-boom value at the 2013 trough, reflecting the collapse in employment net of the rise in productivity per hour. And the maximum loan-to-value ratio is tightened from 100% to 80%. The second is the model's representation of the withdrawal of lenders from high-leverage lending documented by Bover, Torrado, and Villanueva (2019); it is imposed exogenously because the model has a single competitive bank and no mechanism for a lending standard to move on its own.

It is worth being explicit about what the model is not given. There is no shock to beliefs about future

FIGURE 3. The boom and bust: model and data.



NOTES: House prices, consumption and new construction are normalised to their 1997Q4 value. Vertical lines mark the two unanticipated shocks, at 1998Q1 and 2008Q1. Data for prices, consumption and the homeownership rate are quarterly series as in Figure 1; data for construction are the volume index of gross fixed capital formation in dwellings at its 1997 level, its 2007 peak and its 2013 trough.

TABLE 4. House prices, consumption and construction over the boom and bust.

	ΔP_{boom}	ΔC_{boom}	ΔP_{bust}	ΔC_{bust}	I_{peak}/I_{1997}	ΔI_{bust}
Data	123.2%	42.6%	-42.9%	-13.3%	2.16	-52%
Model	113.8%	38.5%	-33.7%	-10.8%	2.16	-31%

NOTES: Boom changes are measured from 1997Q4 to the 2007Q4 peak, bust changes from the 2007Q4 peak to the trough. P is the real house price index, C real aggregate non-durable consumption, I new construction in the model and the volume of gross fixed capital formation in dwellings in the data. The ratio of peak to 1997 construction is the calibration target of Section 4.2; the other five moments are untargeted.

house prices and no shock to the rent-to-price ratio, and construction responds to prices rather than being imposed. Whatever the model generates in house prices, consumption and building, it generates from the series in Figure 1 and the endogenous response of the housing, construction and mortgage markets to them.

5.2. The Model against the Data

Figure 3 and Table 4 compare the model with the data. The model produces a house price boom of 113.8% against 123.2% in the data, and a bust of 33.7% against 42.9%. Consumption rises 38.5% against 42.6% and falls 10.8% against 13.3%: the model undershoots both booms and both busts, the price bust by most. The homeownership rate rises from 79.7% to 81.0% over the boom and falls to 78.9% at the trough, against 79.8%, 80.5% and 77.7% in the data: a little too much entry in the boom and too little exit in the bust. New construction rises by a factor of 2.16 to the peak, which is the one moment of the transition the calibration targets, and falls 31% from the peak to the trough, against 52% in the volume of residential investment between 2007 and 2013: the model has the boom in building and less than the full collapse, because the stock it has accumulated is absorbed by a price fall rather than by years of near-zero construction.


The model undershoots the fall in prices and, by a similar proportion, the fall in consumption. Two further outcomes are worth reporting because the policy experiments act on them directly. Household mortgage debt rises from 0.26 to 0.57 of annual GDP over the boom, against 0.23 to 0.62 in the data: the model has most of the household sector's external borrowing, which is the counterpart of Spain's current account deficit over these years, but not all of it, in keeping with its smaller price boom. And the model generates a foreclosure wave: the foreclosure rate rises from 1.1% before the boom to 2.1% at its bust peak against 1.6% in the data, and the delinquency rate to 4.9% against a peak non-performing loan ratio on household mortgages of 6.3%: the model has fewer delinquencies than the data and more repossessions, because it has no renegotiation, so every delinquent loan faces the same repossession probability each quarter. At the trough, 18.4% of mortgaged households are in negative equity.

Construction deserves a word on its own, because the model and the national accounts measure it differently. At constant prices, the resources the model devotes to building rise from 4.6% of GDP before the boom to 7.6% at the peak, against 6.7% and 10.2% in the chain-linked national accounts. The chain-linked series is at 2010 relative prices, which is why its 1997 share is 6.7% rather than the 6.2% at 1997 prices that Table 3 uses: the level is compared at the prices of the day, the rise in volumes. At current prices the model's share reaches 16.2% of GDP, well above the 11.7% recorded in the data, and the reason is instructive: the model has no land, so the cost of building a house moves one for one with the price of the house, whereas in Spain most of the price rise accrued to land and the cost of construction rose far less. The stock of housing per household ends the boom 10.0% above its starting point and keeps rising through the bust, because building stays above depreciation until prices have fallen a long way; the censuses record a fall in dwellings per household, because the households that immigration created absorbed most of what was built, which is demand the model does not contain.

5.3. What Drove the Boom

Because the drivers enter as separate series, each can be given to the model on its own. Figure 4 shows the boom and bust with only the interest-rate paths, with only the labour-income path, and with both. The rate paths alone raise house prices 72.0% and consumption 22.5% over the boom; the income path alone raises them 30.0% and 13.0%. The two sum to less than the benchmark, for prices and for consumption, because cheaper credit raises the demand for housing by more when incomes are also rising, and the response to either driver is larger when the other is present. Falling rates are the larger of the two forces behind the price boom, and the one that households could not have borrowed their way around: the fall in the mortgage rate raised the house a given payment could buy by more than the rise in income raised the payment a household could make. This decomposition matters for the policy exercise because each driver moves a different denominator, and Section 6.6 uses it to ask how the ranking of the instruments depends on what drives the boom.

FIGURE 4. The drivers one at a time.



NOTES: Real house prices and real consumption, 1997Q4 = 1, when the model is given only the interest-rate paths, only the labour-income path, or both, with the 2008 tightening of lending standards in every case. Vertical lines mark 1998Q1 and 2008Q1.


5.4. Why Consumption Falls

Consumption falls in the model for three distinct reasons, and separating them is what makes the policy evaluation interpretable. The first is the pure *wealth* effect: houses are worth less, and owners who are not constrained in any way consume a little less out of a smaller stock of wealth. The second is the *collateral* effect: falling prices tighten the borrowing constraint (7) and, for households far enough underwater, make the listing constraint $p_t^{\text{ask}} \geq m_t - \tilde{y}_t$ infeasible, so that home equity can no longer be tapped. The third is the *liquidity* effect: even a household with positive equity may be unable to sell within a reasonable horizon, because the market has become slow.

Figure 5 shows the third at work. Expected time on the market for a seller falls from 4.8 months before the boom to 2.9 months at the peak, and then rises to 11.6 months at the trough; the share of the owner-occupied stock listed for sale, 3.1% before the boom, falls to 2.2% at the peak and rises to 7.4% at the trough, which is the overhang the model carries into the bust. Nothing in the model's parameters changes to produce this. It follows from the free-entry conditions: when many households need to sell and few wish to buy, brokers require a larger discount to enter the selling side and a smaller premium on the buying side, and the trading probability at any given price collapses. At the trough the probability that a seller posting the pre-boom discount trades within the quarter is 26%.

To apportion the fall in consumption between the three, I hold each channel fixed in turn over the bust: freezing the price index in households' budget sets at its peak value isolates the wealth effect; holding the borrowing and listing constraints at their peak values isolates the collateral effect; and holding the trading probabilities at their peak values isolates the liquidity effect. Liquidity accounts for 47% of the fall in consumption, the collateral channel for 34%, and the pure wealth effect for 19%. The ordering matches the empirical evidence in Aladangady (2017) and Aruoba, Elul, and Kalemli-Özcan (2022),

FIGURE 5. Housing liquidity over the cycle.




NOTES: Left: the expected time a seller spends on the market, in months, at the seller's chosen listing price. Right: the share of the owner-occupied stock listed for sale. The counterfactual imposes the debt-to-income cap of Section 6 from 1998Q1. Vertical lines mark 1998Q1 and 2008Q1.

which finds that the consumption response to housing wealth is concentrated among constrained households and close to zero for unconstrained ones.

6. Borrower-Based Macroprudential Policy

The counterfactuals impose a borrower-based cap at 1998Q1, unanticipated, on the same initial distribution as the benchmark, and hold it in place thereafter. Everything else—the paths of labour income, interest rates and inflation, and the tightening of lending standards in 2008—is unchanged. Because every regime starts from the same distribution of households, welfare comparisons are exact: each household can be traced across regimes and asked how much it gains. Welfare is measured behind a veil placed at 1998Q1, the moment the boom begins: for each household in the benchmark distribution, I compute the constant proportional increase in non-durable consumption along the benchmark transition that would leave it indifferent to living through the counterfactual transition instead, and average across households. Because the regimes share an initial distribution, no adjustment for differing starting points is needed, and the calculation includes the entire cost of the transition, the years in which constrained households save for a down payment they would not otherwise have needed. Beside every average I report the share of households whose gain is positive, since an average of this kind nets a transfer from those who own houses in 1998 to those who will buy them.

FIGURE 6. Welfare under the three instruments.




NOTES: Left: the consumption-equivalent welfare gain relative to the benchmark, averaged over the 1998Q1 distribution, against the share of pre-boom originations the cap constrains; hollow markers are the best level of each instrument, and the dotted vertical line the common stringency of Section 6.2. Right: the share of households whose gain is positive; the dotted line marks one half.

6.1. Single Instruments

Each instrument is imposed on its own at six levels, from one that would have constrained about a tenth of pre-boom originations to one that would have constrained three quarters of them. Table A1 in Appendix E reports, for every level, the share of pre-boom originations the cap would have constrained, the boom and the bust in prices and consumption, peak debt, peak foreclosure, construction, the welfare gain and the share of households gaining. Figure 6 plots the welfare gain and the share gaining against initial stringency, so that the three instruments share an axis.

Three patterns hold for every instrument. First, a cap dampens the bust far more than the boom. The best DTI cap delivers a price boom of 106.0% rather than 113.8% but cuts the bust from 33.7% to 20.5%; the LTV and PTI caps show the same asymmetry. The cap does not act on the boom directly; it acts on the distribution of leverage that the boom leaves behind, and it is that distribution that determines how far prices fall when demand contracts. Second, the cost of a cap is borne in the boom and it is a consumption cost, not an ownership cost. Constrained households borrow less from abroad, household debt peaks lower, and since output is given, lower borrowing is one for one lower consumption while the borrowing is being done. The homeownership rate at the peak is within a percentage point of the benchmark under every cap at its best level. Third, the welfare gain is hump-shaped in stringency. Tightening keeps reducing exposure to the bust, but it also keeps postponing homeownership, and with $\nu = 0.13$ housing and non-durable consumption are strong complements, so a household kept in the rental sector—capped at quality h_r —suffers a loss that grows quickly as the constraint tightens. Beyond the best level the second force dominates.

FIGURE 7. The transition under the best debt-to-income and loan-to-value caps.



NOTES: House prices and consumption are normalised to their 1997Q4 value in each regime; the third panel is household mortgage debt over annual GDP. The caps are imposed at 1998Q1, unanticipated, on the benchmark initial distribution, at the best level of each instrument. Vertical lines mark 1998Q1 and 2008Q1.

The best levels are a DTI cap of 3.5 times annual earnings, a PTI cap of 25% of earnings, and an LTV cap of 60%. Figure 7 plots the transition under the first and the last against the benchmark. The consumption collapse is halved under the DTI cap, 5.4% against 10.8%, and this does not come mainly from the smaller price fall. It comes from the composition of who is exposed. The foreclosure rate at the bust peak is 1.6% against 2.1%. Under full recourse, a household that is repossessed loses 30% of its income for years; a policy that prevents a household from reaching that state prevents the largest consumption decline any individual household experiences in the model.

6.2. Comparing Instruments

Table 5 compares the three instruments in three ways: at the best level of each, at a matched initial stringency, and in a common currency. All three put the DTI cap first; in a common currency the LTV cap edges the PTI cap on welfare, though not on boom cost.

At the best level of each, the DTI cap delivers the smallest bust and the largest welfare gain, 2.06%, the PTI cap is second with 1.85%, and the LTV cap third with 1.69%; the shares of households gaining rank the instruments the same way. At matched initial stringency, with each cap set to constrain 35% of pre-boom originations so that only the movement of the ceilings over the boom distinguishes them, the order is the same. And in a common currency, setting each cap at the level that reduces the consumption bust to 6.6%, the DTI cap does it at the smallest cost in boom consumption, 2.5% points against 2.7% for the PTI cap and 3.1% for the LTV cap.

Why the instruments differ. A borrower-based cap fixes a ratio: the loan to the value of the house, the loan to the borrower's income, or the loan's payment to the borrower's income. Whatever the cap, the *level* of borrowing it permits is free to move with the denominator, and over the Spanish boom the three denominators moved by very different amounts. Figure 8 makes this explicit on the data. Panel A

TABLE 5. Borrower-based instruments compared.


Regime	Level	Constrained	ΔP_{bust}	ΔC_{bust}	Rise in permitted loan	CEV	Share gaining
Benchmark	—	—	−33.7%	−10.8%	—	—	—
<i>At the best level of each</i>							
DTI	3.5	35%	−20.5%	−5.4%	×1.31	2.06%	58.7%
PTI	25%	47%	−21.3%	−5.9%	×1.44	1.85%	57.5%
LTV	60%	48%	−22.8%	−6.5%	×1.92	1.69%	56.4%
<i>At matched initial stringency (35% of originations constrained)</i>							
DTI	3.5	35%	−20.5%	−5.4%		2.06%	58.7%
PTI	30%	35%	−23.1%	−6.5%		1.58%	56.3%
LTV	70%	35%	−24.5%	−7.1%		1.52%	55.9%
<i>In a common currency (bust in consumption of −6.6%)</i>							
DTI	4.2	24%		−6.6%		1.69%	
PTI	30.5%	34%		−6.6%		1.55%	
LTV	62%	45%		−6.6%		1.65%	

NOTES: Levels: DTI in multiples of annual earnings; PTI as the share of earnings absorbed by the payment on a 25-year amortising loan at the prevailing nominal rate; LTV in per cent. The rise in the permitted loan is the factor by which the maximum loan a cap allows rises between 1997Q4 and 2007Q4 in the model under that cap, to be compared with the same factor in the data in Figure 8 (×2.2 for the LTV cap, ×1.3 for the DTI cap, ×1.4 for the PTI cap). In the common-currency comparison each cap is set at the level that reduces the fall in consumption over the bust to 6.6%; the boom cost in consumption is then 2.5% points for the DTI cap, 2.7% for the PTI cap and 3.1% for the LTV cap.

holds each cap's parameter fixed and asks how much a household of unchanged relative income could borrow against a house of unchanged quality, relative to the end of 1997. The value of the house rose by a factor of 2.2, so an LTV cap permitted more than twice the loan at the peak that it permitted at the start. Labour income rose by a factor of 1.3, so a DTI cap barely moved. The loan that a given share of income could service on a standard amortising loan at the prevailing nominal rate rose by a factor of 1.4 by the end of 2007, and by 1.7 in 2005, when rates were at their lowest, as incomes rose and, for a while, rates fell; a PTI cap assessed at the going rate loosened as rates fell and tightened again as they rose.

Panel B puts the same facts in the terms a buyer faces. Under cap levels chosen so that all three require the same down payment at the end of 1997, the down payment required by an LTV cap rose from 0.7 to 1.2 years of income by 2007; the buyer still brought the same share of the price, but of a price that had more than doubled, and carried the resulting debt into the bust. The down payment required by a DTI cap rose to 3.2 years of income by the end of 2007. The down payment required by a PTI cap at the prevailing rate fell almost to nothing at the rate trough and had climbed back to 2.9 years by the end of 2007. An LTV cap does not loosen over a boom so much as scale with it: it holds leverage constant relative to the collateral and lets it rise relative to income, which is the quantity that decides whether a household can service and refinance its debt once prices fall. A PTI cap indexed to the current rate loosens in response to the shock that drives the boom and, under adjustable rates, re-prices a loan sized at the 2005 rate at the 2008 rate. A DTI cap is tied to the one quantity that barely moved.

FIGURE 8. How far each ceiling rose over the boom.



NOTES: Panel A: the maximum loan a household of unchanged relative income could take against a house of unchanged quality under each cap, with the cap parameter held fixed, 1997Q4 = 1, computed from the data. The LTV ceiling follows the real house price index; the DTI ceiling follows aggregate labour income (employment times productivity per hour); the PTI ceiling follows labour income divided by the payment per unit of principal on a 25-year amortising loan at the prevailing nominal mortgage rate. Panel B: the down payment the same household must bring to buy the same house, in years of income, under cap levels chosen so that all three require the same down payment in 1997Q4 (an LTV cap of 80%, a DTI cap of 2.84 times annual income, a PTI cap of 21% of income); the 1997Q4 price-to-income ratio of 3.5 is the 2002 EFF value scaled back with the price and income indices. Sources: Bank for International Settlements, Eurostat, Banco de España; Appendix A.

The model narrows these differences without removing them, and the column of Table 5 that reports the rise in the permitted loan says something the data alone could not. In the data, the collateral ceiling rose by a factor of 2.2; in the model under the LTV cap it rises by 1.92, because the cap itself holds prices down. General equilibrium therefore narrows the gap between the instruments: a reading of Figure 8 alone would overstate the advantage of the income-based caps. It does not close the gap, because the LTV ceiling still rises more than the PTI ceiling, which in turn rises more than the DTI ceiling, whose denominator the cap cannot move at all since labour income is given. The instruments differ by less than the data suggest and in the order the data predict.

6.3. Payment Caps at a Stressed Rate

The PTI cap's weakness is a weakness of assessing the payment at the current rate. Supervisors who use payment-based limits now assess the payment at a stressed rate rather than the prevailing one, typically the prevailing rate plus a fixed add-on, and that is the practice national authorities have adopted (European Systemic Risk Board 2019). Assessing the best PTI cap's payment at the prevailing rate plus 300 basis points raises its welfare gain from 1.85% to 1.97%, cuts the consumption bust to

TABLE 6. Decomposition of the welfare gain.

Regime	Boom only	Prices held	General equilibrium	Own protection	Externality	Externality share
DTI	−0.58%	0.62%	2.06%	1.20%	1.44%	70%
PTI	−0.52%	0.59%	1.85%	1.11%	1.26%	68%
LTV	−0.49%	0.57%	1.69%	1.06%	1.12%	66%
LTV and DTI	−0.64%	0.64%	2.21%	1.28%	1.57%	71%

NOTES: Consumption-equivalent welfare gain relative to the benchmark, averaged over the 1998Q1 distribution, for each instrument at its best level. Boom only: the cap is imposed and the 2008 shock never arrives. Prices held: the 2008 shock arrives but the house price index follows its benchmark path. General equilibrium: the paper's measure. Own protection is the second column less the first; externality is the third less the second; the share is the externality over the general-equilibrium gain.


5.6%, and places it between the prevailing-rate PTI cap and the DTI cap on every outcome of the bust and on welfare, with a permitted loan that rises by 1.36 rather than 1.44. The reason is in the arithmetic: the add-on damps the response of the permitted principal to the rate, and at the limit, with the payment assessed at a fixed rate, the ceiling rises with income alone and a payment cap is a debt-to-income cap, with the maturity and the assessment rate choosing the multiple. The finding is not that payment-based regulation is unsound, but that a payment-based rule indexed to the prevailing mortgage rate transmits a monetary loosening into a credit loosening, which is precisely the channel Greenwald (2018) identifies, operating here in the direction that defeats the regulation.

6.4. What the Policy Corrects

A cap removes an option, so the welfare gain it produces has to come from somewhere, and Section 3.8 named the source: an externality that runs through the house price. To measure it, I compute the welfare gain from each instrument at its best level in three ways. In the first, the cap is imposed and the boom is permanent: the bust never arrives, and the gain is the pure cost of exclusion, which is negative. In the second, the bust arrives but house prices are held at their benchmark path: the cap protects the households it constrains from the shock they did not anticipate, but no one else. In the third, which is the paper's number, the bust arrives and prices respond: the households the cap constrained were not forced sellers, prices fell less, and every indebted household was better off for it. The difference between the second and the first is the value of the cap's own protection against an unanticipated event, which is what a counterfactual history measures and an optimal-policy exercise would not; the difference between the third and the second is the externality.

Table 6 and Figure 9 report the result. Had the boom been permanent, the DTI cap would have cost the average household 0.58% of lifetime consumption. It is the bust that makes it worth having: with prices held, the cap is worth 0.62%, which is its protection of the households it constrains, 1.20%, net of the cost of exclusion; in general equilibrium the cap is worth 2.06%. Of the total, 70% is the

FIGURE 9. What the policy corrects.



NOTES: The welfare gain from each instrument at its best level, decomposed as in Table 6. The horizontal bar is the general equilibrium total.


smaller price collapse that everyone else enjoys. The pattern is the same for the PTI and LTV caps with smaller magnitudes, and for the combination. The externality is the largest single component of the gain under every instrument, which is the sense in which a cap corrects something rather than merely redistributing: the private cost of a mortgage in this economy is below its social cost by the harm the loan does to other borrowers through the price, and a cap closes part of that gap.

6.5. Who Gains

The average conceals a great deal. Figure 10 plots the distribution of gains under the best DTI cap by initial tenure and Table 7 reports the breakdown; Appendix E gives the same breakdown for the PTI and LTV caps, which is similar. Renters gain most, 5.34%, and 79% of them gain: they buy later, but they buy at lower prices, and they are the group least exposed to the foreclosure wave. Owners with a mortgage gain 3.97%: the cap does not bind on the loan they already hold, but it prevents them from leveraging up further during the boom, and it protects them from the price collapse. Owners without a mortgage lose, 0.09% on average, and only 45% of them gain: they hold the asset whose price the policy holds down and they carry none of the risk it removes. This is the group that would oppose the policy, and it is the largest of the three in Spain, at 53.9% of households, which is why the share of all households gaining, 58.7%, is a majority but not a large one.

The gains are also strongly decreasing in income, 3.86% for the bottom quartile against 0.48% for the top, and among indebted owners they are increasing in initial leverage, 2.78% for the least levered tercile

FIGURE 10. The distribution of welfare gains under the best debt-to-income cap.



NOTES: Distribution across the 1998Q1 population of the consumption-equivalent welfare gain from imposing the DTI cap of 3.5 times annual earnings, by tenure status at 1998Q1. The solid vertical line is zero; the dashed line is the population average.

TABLE 7. Who gains from the best debt-to-income cap.

Group	Share of households	CEV	Share gaining
All households	100%	2.06%	58.7%
<i>By tenure at 1998Q1</i>			
Renters	20.3%	5.34%	79%
Owners with a mortgage	25.8%	3.97%	72%
Owners without a mortgage	53.9%	-0.09%	45%
<i>By income quartile at 1998Q1</i>			
First (lowest)	25%	3.86%	
Second	25%	2.61%	
Third	25%	1.29%	
Fourth (highest)	25%	0.48%	
<i>By initial leverage, indebted owners</i>			
Bottom tercile	8.6%	2.78%	
Middle tercile	8.6%	3.96%	
Top tercile	8.6%	5.19%	

NOTES: Consumption-equivalent welfare gain relative to the benchmark, computed over the whole transition at the 1998Q1 distribution. Groups are defined by the household's state at 1998Q1, before the policy is imposed.

against 5.19% for the most. Borrower-based regulation in this model is, in its incidence, insurance for the levered and the poor paid for by households who own outright.

6.6. Combinations, and What Drives the Boom

Borrower-based regulation is in practice a combination: most European regimes pair a loan-to-value cap with an income-based one. Table 8 adds the combination of the best LTV and DTI caps, and of all

TABLE 8. Combinations, and the ranking under each driver.

Regime	Both drivers (the Spanish boom)				Rates only	Income only
	ΔP_{bust}	ΔC_{bust}	CEV	Share gaining	CEV	CEV
DTI	-20.5%	-5.4%	2.06%	58.7%	1.42%	0.61%
PTI	-21.3%	-5.9%	1.85%	57.5%	1.21%	0.66%
LTV	-22.8%	-6.5%	1.69%	56.4%	0.98%	0.70%
LTV and DTI	-19.4%	-5.0%	2.21%	60.3%	1.49%	0.74%
All three	-19.2%	-4.9%	2.24%	60.6%	1.52%	0.78%

NOTES: Every regime is at the best single levels of Table 5. “Rates only” and “income only” are the booms driven by only the interest-rate and inflation paths or only the labour-income path, as in Figure 4, with the cap imposed at 1998Q1 and the 2008 tightening of lending standards in every case; in the income-only boom rates and inflation stay at their 1997Q4 values, so the payment cap is assessed at a constant rate.

three, to the single instruments. In Spain’s boom the combination adds little: imposing the LTV cap alongside the DTI cap raises the welfare gain from 2.06% to 2.21%, and adding the PTI cap takes it to 2.24%. A combination binds where the tightest cap binds, and in a boom driven by falling rates the income-based cap is the tightest for almost every borrower by the middle of the boom; the collateral cap then constrains only the few households whose income was high relative to the house they wanted.

The value of the combination appears when the boom is driven by something else. The last two columns of Table 8 place each regime in the two one-driver booms of Section 5.3. In the boom driven by rates alone, the ranking is the Spanish one: 1.42% for the DTI cap, 1.21% for the PTI cap, 0.98% for the LTV cap, because that is the boom in which income stands still while prices and payments move. In the boom driven by labour income alone, incomes rise and the DTI ceiling rises with them, the PTI ceiling rises with income too because rates do not move, and prices rise by less than when credit is also cheapening, so the collateral ceiling rises least of all; the three caps end within a few tenths of a point of one another, at 0.61%, 0.66% and 0.70%, and the combination of the LTV and DTI caps, at 0.74%, beats all three, with the three caps together better still at 0.78%. Combining is the best rule in either boom.

What organises these results is that a ratio cap controls debt relative to its denominator, and the cap that works is the one whose denominator moves least over the boom. In Spain that was income, because the boom was one of cheap credit. In a boom of rising incomes it would have been the collateral. An authority that knew the driver could pick the instrument; one that does not should combine, and in the boom that did happen combining cost nothing, since the combination did a little better than the right single instrument and in the other boom it beat every single instrument. The insurance against picking the wrong instrument came free.

7. The Role of Housing Liquidity and Supply

Two features of the economy carry much of what the previous section found: the frictional housing market, and the supply of new housing. This section isolates each.

7.1. Housing Liquidity

The first experiment asks how much of the foregoing depends on the frictional housing market. I replace the search market with a Walrasian one—houses trade at $p_t^h h$ with probability one, brokers and submarkets disappear, and ξ , κ_b , κ_s , γ_b and γ_s drop out—and re-calibrate the remaining parameters to the targets that survive, which are all those of Table 2 except the four that describe housing market activity. The frictionless economy is then given the same shocks and the same caps.

It still produces a boom and a bust, but a much milder consumption cycle: consumption falls 4.9% over the bust against 10.8% in the benchmark, while house prices fall 21.4% against 33.7%. Two things are absent. A household hit by a bad income shock can always sell, so the option to convert housing wealth into consumption never closes, and the delinquency that follows from being unable to sell largely disappears. And the mortgage pricing function no longer embeds expected illiquidity, so credit does not tighten endogenously in the bust. About half of the fall in consumption in the benchmark is the inability to sell.

The consequence for policy evaluation is larger still. The welfare gain from the best DTI cap falls from 2.06% in the benchmark to 0.49% in the frictionless economy, from the best PTI cap from 1.85% to 0.44%, and from the best LTV cap from 1.69% to 0.38%. Roughly three quarters of the value of borrower-based regulation in this model is the value of not being caught holding an asset one cannot sell. The reason goes back to the externality of Section 3.8: when the market is frictionless, a fall in the price is a loss, but a household can always realise what its house is worth and clear its debt if the price covers it; when the market is frictional, the same fall in the price closes the listing constraint for a range of households who then default, and it is the defaults of others that each household's leverage brings about. The friction is what turns a fall in the price into harm done to others, and so what gives a cap most of its value. A model that treats housing as a frictionlessly traded durable will find borrower-based policy to be nearly welfare-neutral, and will do so for a reason that has nothing to do with the policy. Nor is the dependence a cliff: halving the two broker entry costs κ_b and κ_s , holding the other parameters fixed, lowers the gain from the best DTI cap from 2.06% to 1.31% and the expected time on the market at the trough from 11.6 to 7.9 months; the value of regulation falls with the friction smoothly.

7.2. Housing Supply

The second experiment is the construction sector's own. Under the best DTI cap, new construction at the peak of the boom rises by a factor of 1.77 rather than 2.16: the cap removes about a third of the rise in building, because the households it excludes are the marginal buyers whose demand called the building forth. The PTI and LTV caps remove less, 1.87 and 1.92 against 2.16, in the order of their effect on demand. Less building during the boom means a smaller stock to absorb in the bust: the share of the stock listed for sale at the trough, which is the overhang the model carries, is 5.0% under the DTI cap against 7.4% in the benchmark, and expected time on the market at the trough is 8.3 months against 11.6. This is the second way a cap softens the bust, and it is absent from a model with a fixed stock.

The channel runs the other way too. In the frictionless economy, construction at the peak rises by 1.72 rather than 2.16, because a market in which houses can always be sold at the index price generates a smaller price boom for the same credit and income paths, and so calls forth less building; the overhang and the illiquidity reinforce each other, and a model with one and not the other understates the bust. What the supply block cannot do is the part of the Spanish construction boom that was employment: the jobs that building created during the boom and destroyed during the bust are in the model only through the labour-income path it is given. On that side, the model understates both what a cap would have cost in the boom and what it would have saved in the bust.

8. Conclusion

This paper studies whether borrower-based macroprudential policies can mitigate credit-fuelled housing booms and busts and improve welfare. To answer this question, I build and calibrate an equilibrium macro-housing model with heterogeneous households, a frictional housing market, long-term mortgages, and an elastic housing supply. The model captures key features of the Spanish housing boom and bust of 1998–2014, including adjustable-rate, full-recourse mortgages and the absence of borrower-based macroprudential regulation. This setting provides a natural benchmark for evaluating how alternative mortgage lending restrictions would have shaped the cycle.

The results show that borrower-based policies can substantially mitigate the bust in house prices and consumption, although they also dampen the preceding boom. A debt-to-income cap delivers a welfare gain of 2.06% of lifetime consumption for the average household at its welfare-maximising level. Loan-to-value and payment-to-income caps also improve welfare, but their effectiveness differs because each instrument links permitted borrowing to a different underlying variable. In the Spanish baseline, the debt-to-income cap performs best among the individual instruments. Combining a debt-to-income cap with a loan-to-value cap yields an additional welfare gain, albeit a modest one in the baseline, and performs best across alternative scenarios in which the forces driving the boom differ. The results

therefore suggest that the choice of instrument matters, and that combining instruments can provide a more robust response when the source of a future housing boom is uncertain.

The analysis also identifies the mechanisms underlying these welfare gains. First, borrower-based policies address a pecuniary externality that operates through house prices. By limiting leverage, they reduce the number of households that become vulnerable to falling prices and forced sales during a bust. The resulting smaller price decline benefits not only households directly constrained by the policy, but also other indebted households whose ability to sell their homes and repay their mortgages depends on market conditions. This general equilibrium effect is the largest component of the total welfare gain, which is what makes the policy a correction rather than a transfer. Second, housing illiquidity amplifies the effects of the bust. When households cannot sell their homes quickly, adverse income shocks and falling collateral values translate into sharper consumption cuts and a greater risk of default. Removing housing market frictions substantially reduces both the consumption collapse and the welfare value of regulation, highlighting the importance of accounting for the difficulty of selling housing assets when evaluating macroprudential policy. Third, the distributional effects are heterogeneous. Renters and mortgaged homeowners benefit from lower prices and reduced exposure to default, while outright homeowners lose from the lower value of their housing wealth. The aggregate welfare gains therefore mask meaningful differences across households.

These findings have implications for the design of borrower-based macroprudential policy. A loan-to-value cap tied to current house prices can accommodate substantial increases in borrowing during a boom, while a payment-to-income cap tied to prevailing interest rates can loosen when falling rates make a given payment support a larger loan. Debt-to-income limits are less sensitive to these particular movements, which helps explain their performance in the Spanish episode. Payment-to-income limits assessed at stressed interest rates can mitigate this weakness, while combining instruments can reduce reliance on any single measure of borrowing capacity. More broadly, the results suggest that macroprudential authorities should consider not only how restrictive a lending limit is at a given point in time, but also how its denominator evolves over the housing cycle. The welfare case for intervention rests on the external cost of leverage, and it is strongest where housing is illiquid, because illiquidity is what turns a fall in prices into forced sales and defaults.

Several avenues for further research follow from this analysis. First, extending the framework to include banks and their balance-sheet constraints would allow an analysis of the interaction between borrower-based policies, lender behaviour, and financial stability, as well as the potential for credit to migrate towards less regulated intermediaries. Second, studying the optimal time-varying adjustment of lending limits would help establish how authorities should respond to changes in house prices, interest rates, income growth, and credit conditions, rather than setting limits once at the beginning of a boom. Finally, applying the framework to other countries and institutional settings would help determine how the effectiveness of alternative instruments depends on mortgage contract design, recourse laws, housing supply elasticity, and the liquidity of the housing market. These extensions

would help establish how borrower-based macroprudential policy can be designed to remain effective across different housing cycles and financial systems.

References

- Akin, Ozlem, José García Montalvo, Jaume García Villar, José-Luis Peydró, and Josep Maria Raya. 2014. "The Real Estate and Credit Bubble: Evidence from Spain." *SERIEs* 5 (2-3): 223-243.
- Akinci, Ozge and Jane Olmstead-Rumsey. 2018. "How Effective Are Macroprudential Policies? An Empirical Investigation." *Journal of Financial Intermediation* 33: 33-57.
- Aladangady, Aditya. 2017. "Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata." *American Economic Review* 107 (11): 3415-3446.
- Arena, Marco, Tingyun Chen, Seung Mo Choi, Nan Geng, Cheikh A. Gueye, Tonny Lybek, Evan Papageorgiou, and Yuanyan Sophia Zhang. 2020. *Macroprudential Policies and House Prices in Europe*. Washington, DC: International Monetary Fund.
- Aruoba, S. Borağan, Ronel Elul, and Şebnem Kalemli-Özcan. 2022. "Housing Wealth and Consumption: The Role of Heterogeneous Credit Constraints." Working paper, National Bureau of Economic Research.
- Banco de España. 2020. "The Housing Market in Spain: 2014-2019." Documentos Ocasionales 2013, Banco de España.
- Bianchi, Javier. 2011. "Overborrowing and Systemic Externalities in the Business Cycle." *American Economic Review* 101 (7): 3400-3426.
- Bover, Olympia, María Torrado, and Ernesto Villanueva. 2019. "The Loan to Value Ratio for Housing in Spain over the Period 2004-2016." *Banco de España Economic Bulletin* 1/2019: 1-9.
- Cano-Fuentes, Gala, Aitziber Etxezarreta Etxarri, Kees Dol, and Joris Hoekstra. 2013. "From Housing Bubble to Repossessions: Spain Compared to Other West European Countries." *Housing Studies* 28 (8): 1197-1217.
- Carro, Adrián. 2022. "Could Spain Be Less Different? Exploring the Effects of Macroprudential Policy on the House Price Cycle." Documentos de Trabajo 2230, Banco de España.
- Carro, Adrián, Marc Hinterschweiger, Arzu Uluc, and J. Doyne Farmer. 2022. "Heterogeneous Effects and Spillovers of Macroprudential Policy in an Agent-Based Model of the UK Housing Market." Documentos de Trabajo 2217, Banco de España.
- Cerutti, Eugenio, Stijn Claessens, and Luc Laeven. 2017. "The Use and Effectiveness of Macroprudential Policies: New Evidence." *Journal of Financial Stability* 28: 203-224.
- Cloyne, James, Clodomiro Ferreira, and Paolo Surico. 2020. "Monetary Policy when Households Have Debt: New Evidence on the Transmission Mechanism." *Review of Economic Studies* 87 (1): 102-129.
- Congreso de los Diputados. 2009. "Comparecencia de la Ministra de Vivienda ante la Comisión de Vivienda." Diario de Sesiones del Congreso de los Diputados, Comisiones, núm. 449, IX Legislatura, 16 de diciembre de 2009.
- Corbae, Dean and Erwan Quintin. 2015. "Leverage and the Foreclosure Crisis." *Journal of Political Economy* 123 (1): 1-65.
- Dávila, Eduardo and Anton Korinek. 2018. "Pecuniary Externalities in Economies with Financial Frictions." *Review of Economic Studies* 85 (1): 352-395.
- Estrada García, Ángel and Jesús Saurina Salas. 2016. "Spanish Boom-Bust and Macroprudential Policy." *Banco de España Financial Stability Review* 30: 35-61.

- European Systemic Risk Board. 2019. “A Review of Macroprudential Policy in the EU in 2018.” , European Systemic Risk Board, Frankfurt.
- 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.
- Fernández-Villaverde, Jesús, Luis Garicano, and Tano Santos. 2013. “Political Credit Cycles: The Case of the Eurozone.” *Journal of Economic Perspectives* 27 (3): 145–166.
- Ferreira, Clodomiro, Julio Gálvez, and Myroslav Pidkuyko. 2024. “Housing Tenure, Consumption and Household Debt: Life-Cycle Dynamics during a Housing Bust in Spain.” Documentos de Trabajo 2424, Banco de España.
- Flavin, Marjorie and Shinobu Nakagawa. 2008. “A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence.” *American Economic Review* 98 (1): 474–495.
- Fuentes Castro, Daniel. 2009. “A note on housing supply excess and the duration of the sector adjustment.” MPRA Paper 51609, University Library of Munich.
- Galán, Jorge E. and Matías Lamas. 2019. “Beyond the LTV Ratio: New Macroprudential Lessons from Spain.” Documentos de Trabajo 1931, Banco de España.
- Galán, Jorge E., Matías Lamas, and Raquel Vegas. 2022. “Roots and Recourse Mortgages: Handing Back the Keys.” Documentos de Trabajo 2203, Banco de España.
- García-Montalvo, José and Josep Maria Raya. 2018. “Constraints on LTV as a Macroprudential Tool: A Precautionary Tale.” *Oxford Economic Papers* 70 (3): 821–845.
- 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, Rodolfo Manuelli, and Adrian Peralta-Alva. 2019. “A Macroeconomic Model of Price Swings in the Housing Market.” *American Economic Review* 109 (6): 2036–2072.
- Ghirelli, Corinna, Danilo Leiva-León, and Alberto Urtasun. 2022. “Housing Prices in Spain: Convergence or Decoupling?” Documentos de Trabajo 2205, Banco de España.
- Greenwald, Daniel L. 2018. “The Mortgage Credit Channel of Macroeconomic Transmission.” Working Paper 5184-16, MIT Sloan School of Management.
- Grodecka, Anna. 2020. “On the Effectiveness of Loan-to-Value Regulation in a Multiconstraint Framework.” *Journal of Money, Credit and Banking* 52 (5): 1231–1270.
- Guner, Nezih, Javier López-Segovia, and Roberto Ramos. 2020. “Reforming the Individual Income Tax in Spain.” *SERIEs* 11 (4): 369–406.
- Guren, Adam M., Alisdair McKay, Emi Nakamura, and Jón Steinsson. 2021. “Housing Wealth Effects: The Long View.” *Review of Economic Studies* 88 (2): 669–707.
- Hedlund, Aaron. 2016. “The Cyclical Dynamics of Illiquid Housing, Debt, and Foreclosures.” *Quantitative Economics* 7 (1): 289–328.
- International Monetary Fund. 2014. “Staff Guidance Note on Macroprudential Policy.” , International Monetary Fund, Washington, DC.
- Jimeno, Juan F. and Tano Santos. 2014. “The Crisis of the Spanish Economy.” *SERIEs* 5 (2–3): 125–141.
- 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.
- Kuttner, Kenneth N. and Ilhyock Shim. 2016. “Can Non-Interest Rate Policies Stabilize Housing Markets? Evidence

- from a Panel of 57 Economies.” *Journal of Financial Stability* 26: 31–44.
- Laliotis, Dimitrios, Alejandro Buesa, Miha Leber, and Javier Población. 2020. “An Agent-Based Model for the Assessment of LTV Caps.” *Quantitative Finance* 20 (10): 1721–1748.
- Lorenzoni, Guido. 2008. “Inefficient Credit Booms.” *Review of Economic Studies* 75 (3): 809–833.
- Martín, Alberto, Enrique Moral-Benito, and Tom Schmitz. 2021. “The Financial Transmission of Housing Booms: Evidence from Spain.” *American Economic Review* 111 (3): 1013–1053.
- Mas Ivars, Matilde and Francisco Pérez García. 2021. “El stock de capital en España y sus comunidades autónomas. Revisión metodológica y evolución reciente de la inversión y el capital 1995–2020.” Documento de Trabajo 1/2021, Fundación BBVA–Ivie.
- Menzio, Guido and Shouyong Shi. 2010. “Block Recursive Equilibria for Stochastic Models of Search on the Job.” *Journal of Economic Theory* 145 (4): 1453–1494.
- Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. “Household Balance Sheets, Consumption, and the Economic Slump.” *Quarterly Journal of Economics* 128 (4): 1687–1726.
- Mian, Atif and Amir Sufi. 2011. “House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis.” *American Economic Review* 101 (5): 2132–2156.
- Mian, Atif and Amir Sufi. 2014. “What Explains the 2007–2009 Drop in Employment?” *Econometrica* 82 (6): 2197–2223.
- Moen, Espen R. 1997. “Competitive Search Equilibrium.” *Journal of Political Economy* 105 (2): 385–411.
- Rouwenhorst, K. Geert. 1995. “Asset Pricing Implications of Equilibrium Business Cycle Models.” *Frontiers of Business Cycle Research* pp. 294–330.
- Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron. 2004. “Consumption and Risk Sharing over the Life Cycle.” *Journal of Monetary Economics* 51 (3): 609–633.

Appendix A. Data

Dating. The pre-boom reference point is 1997Q4, the quarter the model’s initial stationary equilibrium represents. Boom changes are measured from 1997Q4 to 2007Q4 and bust changes from 2007Q4 to the trough of each series, which for house prices falls in 2013. The comparison of the three borrower-based caps in Section 6.2 uses the same dates.

House prices. Real residential property prices for Spain and the euro area are the Bank for International Settlements selected series, quarterly, deflated by the consumer price index. In the text they are normalised to 1997Q4; in Figure 1, to their 1996–1997 average.

Consumption. Real private final consumption expenditure for Spain, quarterly, seasonally adjusted, in chained volumes, from the OECD national accounts (FRED series NAEXKP02ESQ189S).

Employment and productivity. The employment rate of the population aged 20 to 64, quarterly and seasonally adjusted, and real labour productivity per hour worked, quarterly, both from Eurostat; the comparison across euro-area members in Section 2.2 uses the same source. Aggregate labour income

per adult is the product of the two. The model input is the product of the two, which the text calls aggregate labour income.

Interest rates and inflation. The nominal mortgage rate is the Banco de España reference series for the mortgage market, the quarterly average rate on mortgage credit (series D_1T9H0001); the nominal risk-free rate is the one-year Treasury bill rate from the same source. Inflation is the year-on-year rate of the Spanish consumer price index, monthly, from the OECD Main Economic Indicators (FRED series CPALTT01ESM659N), averaged to quarters; the Spanish HICP, which starts in 1997, gives the same figures to a tenth of a point. Real rates are obtained by exact deflation, $(1 + i)/(1 + \pi) - 1$, with the same inflation series and timing throughout; annual rates are converted to quarterly once, as $(1 + i)^{1/4} - 1$. The pre-boom values used in the calibration are 1996–1997 averages; the boom and bust values quoted in the text are averages over 1998Q1–2007Q4 and 2008Q1–2014Q4.

Construction. Residential investment is gross fixed capital formation in dwellings from the annual national accounts, compiled by INE and disseminated by Eurostat (table nama_10_an6, asset N111), at current prices and in chain-linked volumes at 2010 prices; shares of GDP use GDP from the same accounts on the same basis (table nama_10_gdp). The volume index peaks in 2007 and reaches its trough in 2013; both shares of GDP peak in 2006. Housing starts are the Ministerio de Vivienda quarterly series of dwellings started (*viviendas iniciadas*), summed to calendar years, as reproduced in the annex of Fuentes Castro (2009). The stock of newly built dwellings without a buyer is the Ministerio de Vivienda count of completed new dwellings for sale at 31 December 2008, reported to the Congreso de los Diputados on 16 December 2009 (Congreso de los Diputados 2009). Dwellings and households are from the INE population and housing censuses of 2001 and 2011. The depreciation rate, 2.0% a year, is set above the geometric rate for dwellings in the current Fundación BBVA–Ivie capital stock series, 1.14% (Mas Ivars and Pérez García 2021), for the reason given in Section 4.

The three borrower-based caps. Panel A of Figure 8 divides each series by its 1997Q4 value: the real house price index for the LTV ceiling; aggregate labour income for the DTI ceiling; and aggregate labour income divided by the payment per unit of principal on a 100-quarter amortising loan at the quarterly nominal mortgage rate for the PTI ceiling. Two variants are reported for reference and not plotted: with an interest-only payment at the nominal rate the PTI factor at 2007Q4 is 1.6; with an interest-only payment at the real rate it is undefined, because the real rate falls to -1.22% in 2006. Panel B needs a price-to-income level: the 2002 EFF ratio of housing wealth to mean annual earnings, 4.37, scaled back to 1997Q4 with the price and labour-income indices, gives 3.5. The cap levels in Panel B are chosen so that all three require the same down payment in 1997Q4: an LTV cap of 80%, and the DTI and PTI levels that permit the same loan, 2.84 times annual income and 21% of income respectively.

Household debt. Household mortgage debt over annual GDP is lending to households for house purchase from the Banco de España Financial Accounts, over nominal GDP from INE.

Loan-to-value ratios. The distribution of loan-to-value ratios at origination is from Bover, Torrado, and Villanueva (2019), who match the universe of mortgage deeds to appraisal values.

Household portfolios. Distributional targets are computed from the 2002 wave of the *Encuesta Financiera de las Familias*, the earliest available, using the first impute and the survey weights. House values are self-reported current values; loan-to-value ratios at origination are the ratio of the initial principal to the purchase price, trimmed at 1.5 times the interquartile range.

Homeownership. The homeownership rate is the INE series on tenure of main dwelling, interpolated to quarterly frequency between census and survey observations.

Foreclosures. Foreclosure filings and the duration of the judicial foreclosure procedure are from the statistics of the General Council of the Judiciary (CGPJ). The delinquency rate in the data is the Banco de España non-performing loan ratio on lending to households for house purchase.

Appendix B. The Mortgage Pricing Function

This appendix states the objects that enter (12). Throughout, expectations are over (z_{t+1}, e_{t+1}) conditional on z_t , and starred variables are the borrower's date- $(t+1)$ policies at the state $(y_{t+1}, z_{t+1}, m_{t+1}, h_t)$.

Continuation after a payment. A borrower who pays l_{t+1}^* carries the balance $m_{t+2}^* = (m_{t+1} - l_{t+1}^*)(1 + r_{t+1}^m)$ into the following period. The bank marks this to market at $q_{t+1}^c \equiv q_{t+1}(y_{t+2}^*, z_{t+1}, h_t, m_{t+2}^*)$, that is, at the price it would charge the same borrower for the same balance were the loan originated afresh. The recursion is therefore in the same function q , and is solved jointly with the household's problem.

Continuation after a skipped payment. A borrower who is delinquent and is not repossessed carries the balance forward, increased by any maintenance the bank advanced, $m_{t+1} + a_{t+1}$ with $a_{t+1} = \max\{0, v p_{t+1}^h h_t - y_{t+1}\}$. The bank marks that balance at $q_{t+1}^s \equiv q_{t+1}(y_{t+2}^*, z_{t+1}, h_t, m_{t+1} + a_{t+1})$ evaluated at the policies of a household that skipped, which differ from those of a paying household because its cash on hand is higher by the foregone payment, and its value in this branch is the marked balance less the advance it paid out.

The deficiency judgment. If the bank repossesses, it sells the house subject to the same market frictions faced by any seller and, in expectation over the listing outcome, recovers $(1 - \chi) p_{t+1}^h h_t - M_{t+1}$, where M_{t+1} is the maintenance it expects to pay while it holds the house, $v p^h h_t$ per quarter over the expected time to sale. The borrower owes $d_{t+1} = \max\{0, m_{t+1} - (1 - \chi) p_{t+1}^h h_t\}$ on the balance including any earlier

advance, of which the bank recovers the discounted garnishment stream

$$Q_{t+1}^{\text{dj}} = \mathbb{E}_{t+1} \sum_{j \geq 0} \left(\prod_{i=0}^j \frac{1}{1 + r_{t+1+i}^{\text{m}}} \right) \min \left\{ \tau y_{t+1+j}, d_{t+1+j} \right\},$$

where d_{t+1+j} follows the law of motion in Section 3.8 and the household's income path is generated by its policies under V^{d} . Since the garnishment is a fraction of income and income is bounded, $Q_{t+1}^{\text{dj}} \leq d_{t+1}$, with equality only in the limit of a household that never receives a low draw.

Branch indicators. $\mathbf{1}_{t+1}^{\text{m}*}$, $\mathbf{1}_{t+1}^{\text{p}*}$ and $\mathbf{1}_{t+1}^{\text{d}*}$ are the indicators of which of the three branches attains the maximum in W_{t+1}^{h} , evaluated at the borrower's optimal policies. They depend on q_{t+1} through the refinancing branch, so the pricing function is determined jointly with the household's problem; Appendix C describes how.

Appendix C. Computation

Stationary equilibrium. The state space is discretised on a grid for cash on hand, a grid for the mortgage balance expressed as pseudo-leverage, the four productivity states, and the three house qualities, with a finer auxiliary grid used to compute moments. The stationary equilibrium is a fixed point in the price index p^{h} , computed by bisection. For a candidate p^{h} :

- a. Construct the tightness schedules (4) and the implied trading probabilities on the grid of posted prices, with the selling probability capped at one.
- b. Solve the household's problem and the pricing function jointly by backward iteration, alternating a value function update given q with a pricing update given the policies, until both converge. Policy iteration is used to accelerate the value function step once the policies stop changing. The maintenance test of Section 3.8 is applied in the listing sub-period before the listing choice is evaluated.
- c. Iterate the distribution to its invariant point under the converged policies, splitting the update into the listing, mortgage, and bidding sub-periods.
- d. Compute excess demand for housing quality—what brokers sell to buyers and to owners and the bank for maintenance, less what they buy from sellers, from reposessions and from builders at the construction level (5) evaluated at the candidate price—and update p^{h} . In the stationary equilibrium construction equals depreciation, so the stock is whatever the candidate price and the distribution imply, and the housing-wealth target of Table 2 pins it.

Transition. The transition is solved by a shooting algorithm on the sequence $\{p_t^{\text{h}}\}_{t=0}^T$, with T set so that the economy is within tolerance of the terminal stationary equilibrium. Given a candidate sequence, construction and the stock follow from (5) and (6) directly, the tightness schedules follow from (4),

value functions and pricing functions are solved backward from the terminal steady state, and the distribution is iterated forward from Φ_0 ; the excess demand at each date, which now includes net additions to the stock, updates the sequence. The terminal stationary equilibrium is the one whose stock the path leads to, so it is solved jointly with the path by iterating on its stock. An unanticipated shock at date s is implemented by solving the segment $[s, T]$ as its own transition, with the terminal stationary equilibrium of the new regime and with Φ_s and H_s inherited from the previous segment. The boom–bust experiment therefore consists of two such segments, at 1998Q1 and 2008Q1. The nominal rate that enters the payment cap is the observed series until the data end and the terminal real rate plus two per cent inflation thereafter.

Counterfactual regimes. A single-instrument regime replaces the benchmark’s origination constraint by the cap at the stated level from 1998Q1 and solves both segments afresh from the benchmark’s 1998Q1 distribution and stock. The share of pre-boom originations a cap constrains is computed in the benchmark stationary equilibrium by applying the cap to the observed originations. Combinations impose the caps jointly. The common-currency levels of Section 6.2 are found by bisection on the cap level to the stated bust in consumption. The driver regimes of Section 6.6 feed only the interest-rate and inflation paths, or only the labour-income path, holding the other at its 1997Q4 value through both segments, with the 2008 tightening of lending standards in every case. The stressed-rate payment cap adds a constant to the nominal rate at which the payment in (9) is assessed.

Welfare. The consumption-equivalent variation for a household at state x in the 1998Q1 distribution solves $V_{1998}^{\text{bench}}(x; 1 + g(x)) = V_{1998}^{\text{cf}}(x)$, where the left-hand side is the benchmark value with non-durable consumption scaled by $1 + g$ at every date and state along the transition. Because u is homothetic in (c, c_h) but the scaling applies only to c , $g(x)$ is computed by univariate root-finding at each state rather than in closed form. The aggregate figure reported in the text is $\int g(x) d\Phi_{1998}(x)$, and the share gaining is $\int \mathbf{1}\{g(x) > 0\} d\Phi_{1998}(x)$. The three runs of Section 6.4 are: the cap imposed with the first segment solved to its own terminal stationary equilibrium and no second segment; the cap imposed with both segments but with the household’s problem solved along the benchmark’s price sequence, which requires no fixed point; and the full general-equilibrium transition.

Frictionless economy. The Walrasian economy of Section 7.1 replaces the bidding and listing sub-periods by a tenure choice at the index price, drops the four search parameters and the cost of a failed listing, and re-solves the calibration of Section 4.2 for the remaining parameters to the targets that survive. Its stationary equilibrium and transitions are computed as above with the trading probabilities set to one.

Appendix D. An Alternative Earnings Process

The benchmark takes the earnings process from Guner, López-Segovia, and Ramos (2020). This appendix reports an alternative estimated directly on the 2004–2012 Spanish EU-SILC, and shows that the results are not sensitive to the choice.

Following Storesletten, Telmer, and Yaron (2004), log earnings of individual i at age h are

$$\begin{aligned} w_{ih} &= \alpha_i + z_{ih} + \varepsilon_{ih}, \\ z_{ih} &= \rho z_{ih-1} + \eta_{ih}, \end{aligned}$$

where $h = 0, \dots, 39$ normalises a working life from 25 to 64, $\alpha_i \sim \mathcal{N}(0, \sigma_\alpha^2)$ is a fixed effect, $\eta_{ih} \sim \mathcal{N}(0, \sigma_\eta^2)$ and $\varepsilon_{ih} \sim \mathcal{N}(0, \sigma_\varepsilon^2)$, all mutually independent. The variance of log earnings at age h is

$$\mathbb{V}_i[w_{ih}] = \sigma_\varepsilon^2 + \sigma_\alpha^2 + \sigma_\eta^2 \sum_{j=0}^{h-1} \rho^{2j}, \quad (\text{A1})$$

and the autocovariance at lag n is $\mathbb{C}_i[w_{ih+n}, w_{ih}] = \sigma_\alpha^2 + \rho^n \sigma_\eta^2 / (1 - \rho^2)$. The identifying margins are the variance at labour market entry, $\mathbb{V}_i[w_{i0}] = \sigma_\varepsilon^2 + \sigma_\alpha^2$; its growth with age, $\Delta_h \mathbb{V}_i[w_{ih}] = \sigma_\eta^2 \rho^{2h}$; the curvature of that growth, $\Delta_h^2 \mathbb{V}_i[w_{ih}] / \Delta_h \mathbb{V}_i[w_{ih}] = 1 + \rho^{-2}$; and the autocovariances, which separate the fixed effect from the persistent component.

Minimising the distance between the variance and autocovariance profiles in the data and in a simulated sample of 20,000 individuals drawn from (A1) yields $\rho = 0.9890$, $\sigma_\alpha = 0.4677$, $\sigma_\eta = 0.1732$ and $\sigma_\varepsilon = 0.2249$. These are annual values; converted to the quarterly model period as in Section 4.1, the persistence is 0.9972. Discretising as in Section 4.1 and re-calibrating the internal parameters to the same targets leaves the transition results in Table 4 unchanged to within half a percentage point, and the ranking of the three instruments in Section 6 unchanged. The reason the choice matters so little is that the housing decisions in the model are driven by the persistence of the income process and by the size of the upper tail, and the two specifications agree closely on both.

Appendix E. Additional Tables

Table A1 reports the full set of outcomes for the single-instrument experiments of Section 6.1. Table A2 reports the incidence of the best payment-to-income and loan-to-value caps, as Table 7 does for the debt-to-income cap. Table A3 reports the benchmark boom and bust under each driver on its own.

TABLE A1. Single instruments at six levels: all outcomes.

Cap	Constr.	ΔP_{boom}	ΔC_{boom}	ΔP_{bust}	ΔC_{bust}	Debt/GDP	Foreclosure	I_{peak}/I_{1997}	CEV	Gaining
Benchmark	—	113.8%	38.5%	−33.7%	−10.8%	0.57	2.1%	2.16	—	—
<i>Loan-to-value cap, ϑ</i>										
90%	12%	111.7%	37.6%	−29.5%	−9.0%	0.52	1.9%	2.07	0.31%	38.4%
80%	22%	110.2%	36.9%	−26.9%	−8.0%	0.49	1.9%	2.02	0.94%	51.7%
70%	35%	108.4%	36.1%	−24.5%	−7.1%	0.46	1.8%	1.96	1.52%	55.9%
60%	48%	106.6%	35.2%	−22.8%	−6.5%	0.44	1.7%	1.92	1.69%	56.4%
50%	62%	104.7%	34.4%	−21.6%	−6.0%	0.42	1.7%	1.89	1.41%	52.4%
40%	76%	102.9%	33.6%	−20.7%	−5.8%	0.41	1.6%	1.87	0.62%	44.1%
<i>Debt-to-income cap, δ (times annual earnings)</i>										
5.0	10%	111.3%	37.4%	−28.5%	−8.6%	0.51	1.9%	2.01	0.72%	42.0%
4.5	19%	109.3%	36.5%	−24.9%	−7.2%	0.47	1.8%	1.91	1.38%	52.6%
4.0	27%	107.6%	35.7%	−22.5%	−6.2%	0.44	1.7%	1.83	1.84%	57.1%
3.5	35%	106.0%	35.0%	−20.5%	−5.4%	0.41	1.6%	1.77	2.06%	58.7%
3.0	46%	103.9%	34.0%	−18.4%	−4.7%	0.39	1.6%	1.70	1.77%	55.2%
2.5	60%	101.2%	32.8%	−16.6%	−4.1%	0.36	1.5%	1.64	0.98%	46.9%
<i>Payment-to-income cap, ψ (share of earnings)</i>										
40%	14%	111.1%	37.3%	−28.2%	−8.5%	0.51	1.9%	2.04	0.44%	39.3%
35%	24%	109.4%	36.5%	−25.4%	−7.4%	0.47	1.8%	1.97	1.02%	50.9%
30%	35%	107.6%	35.7%	−23.1%	−6.5%	0.44	1.7%	1.92	1.58%	56.3%
25%	47%	105.7%	34.9%	−21.3%	−5.9%	0.42	1.7%	1.87	1.85%	57.5%
20%	61%	103.5%	33.9%	−19.9%	−5.4%	0.40	1.6%	1.83	1.62%	54.0%
15%	75%	101.5%	32.9%	−18.9%	−5.1%	0.39	1.6%	1.80	0.87%	45.6%
<i>Combinations and the stressed-rate payment cap</i>										
LTV and DTI	52%	104.9%	34.7%	−19.4%	−5.0%	0.40	1.6%	1.73	2.21%	60.3%
All three	58%	104.7%	34.6%	−19.2%	−4.9%	0.39	1.5%	1.72	2.24%	60.6%
PTI, stressed	49%	105.9%	34.8%	−20.9%	−5.6%	0.42	1.6%	1.81	1.97%	58.1%

NOTES: Each row imposes the stated cap at 1998Q1 on the benchmark initial distribution and holds it thereafter; all other inputs are as in Section 5.1. “Constr.” is the share of originations in the pre-boom stationary equilibrium that the cap would have bound; CEV is the consumption-equivalent welfare gain relative to the benchmark, averaged over the 1998Q1 distribution; “Gaining” is the share of those households whose gain is positive. Debt/GDP is household mortgage debt over annual GDP at the 2007Q4 peak; foreclosure is the maximum quarterly foreclosure rate over the bust, annualised; I_{peak}/I_{1997} is new construction at its peak relative to 1997Q4. Combinations are at the best single levels; the stressed payment cap assesses the payment at the prevailing rate plus 300 basis points.

TABLE A2. Who gains from the best payment-to-income and loan-to-value caps.

Group	Share of households	PTI cap, $\psi = 25\%$		LTV cap, $\vartheta = 60\%$	
		CEV	Share gaining	CEV	Share gaining
All households	100%	1.85%	57.5%	1.69%	56.4%
<i>By tenure at 1998Q1</i>					
Renters	20.3%	4.81%	78%	4.43%	76%
Owners with a mortgage	25.8%	3.58%	70%	3.29%	69%
Owners without a mortgage	53.9%	-0.09%	44%	-0.11%	43%
<i>By income quartile at 1998Q1</i>					
First (lowest)	25%	3.47%		3.17%	
Second	25%	2.34%		2.14%	
Third	25%	1.16%		1.06%	
Fourth (highest)	25%	0.43%		0.39%	
<i>By initial leverage, indebted owners</i>					
Bottom tercile	8.6%	2.52%		2.31%	
Middle tercile	8.6%	3.57%		3.28%	
Top tercile	8.6%	4.67%		4.31%	

NOTES: As Table 7.

TABLE A3. The benchmark boom and bust under each driver.

Drivers	ΔP_{boom}	ΔC_{boom}	ΔP_{bust}	ΔC_{bust}
Both (benchmark)	113.8%	38.5%	-33.7%	-10.8%
Interest rates and inflation only	72.0%	22.5%	-23.1%	-6.9%
Labour income only	30.0%	13.0%	-11.2%	-3.7%

NOTES: As Figure 4. The 2008 tightening of lending standards is imposed in every case.