

Laliotis, Dimitrios; Buesa, Alejandro; Leber, Miha; Población García, Javier

Working Paper

An agent-based model for the assessment of LTV caps

ECB Working Paper, No. 2294

Provided in Cooperation with:

European Central Bank (ECB)

Suggested Citation: Laliotis, Dimitrios; Buesa, Alejandro; Leber, Miha; Población García, Javier (2019) : An agent-based model for the assessment of LTV caps, ECB Working Paper, No. 2294, ISBN 978-92-899-3556-2, European Central Bank (ECB), Frankfurt a. M., <https://doi.org/10.2866/71742>

This Version is available at:

<https://hdl.handle.net/10419/208328>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



EUROPEAN CENTRAL BANK

EUROSYSTEM

Working Paper Series

Dimitrios Laliotis, Alejandro Buesa,
Miha Leber and
Francisco Javier Población García

An agent-based model for the
assessment of LTV caps

No 2294 / July 2019

Abstract

We assess the effects of regulatory caps in the loan-to-value (LTV) ratio using agent-based models (ABMs). Our approach builds upon a straightforward ABM where we model the interactions of sellers, buyers and banks within a computational framework that enables the application of LTV caps. The results are first presented using simulated data and then we calibrate the probability distributions based on actual European data from the HFCS survey. The results suggest that this approach can be viewed as a useful alternative to the existing analytical frameworks for assessing the impact of macroprudential measures, mainly due to the very few assumptions the method relies upon and the ability to easily incorporate additional and more complex features related to the behavioral response of borrowers to such measures.

Keywords: Borrower-based measures, macroprudential policy, house prices, HFCS survey.

JEL codes: D14, D31, E50, R21.

Non-technical summary

Since housing is one of the most important sectors of the economy, the development of real estate bubbles and crashes plays a very important role in almost all financial crises. The complexity and interrelation of housing market characteristics makes the modelling of housing sector-related macroprudential measures often highly challenging.

In this paper, we assess LTV cap measures in the housing market using an agent-based model (ABM) which consists of home property sellers, buyers and banks. The economic environment, in the form of credit provision by banks and the applied LTV caps, is mapped onto behavioral reactions by the different agent types that eventually result in the emergence of a one-step housing market settlement/clearance state, where specific properties are sold at specific prices. The market-clearing mechanism consists of parallel auctions of different quality properties, allowing for different types of seller and buyer agents. Each buyer agent will always go for the highest-quality property first but, with a given probability, he/she will also decide to participate in auctions of lower quality and lower starting asking price properties. With another given probability, they will opt for a higher down payment to overcome potential credit restrictions imposed by the application of an LTV cap measure.

Results based on simulated data are presented first. In this case we assume that initial liquid wealth, total wealth, LTV at origination and property value parameters follow three types of probability distributions in which initial liquid wealth and LTV at origination can be correlated (positively or negatively). Since the cap constraint becomes binding for a significant proportion of households under the assumption that there is no change in their household's liquid assets and their ability to come up with the required down payment, the application of an LTV cap naturally shifts the distribution of buyers towards lower price ranges.

After this simulation exercise, the relevant probability distributions are calibrated on actual European data. In that context, the second wave of the Household Finance and Consumption Survey (HFCS) is used. We also deploy a copula-based approach for the estimation of multivariate distributions of initial liquid wealth, debt-service-to-income ratios, LTV ratios and HMR Value at origination.

We find that in this database high LTVs correspond mainly to the low and medium value of the house price range. Hence, when we impose a cap on LTV, low and medium value houses are more affected than high value houses. This evidence reveals the presence of "sub-prime" mortgage loan segments in the European loan markets, which is one potentially significant feature with macroprudential policy implications.

When impacts are calculated country by country, it is understood that these results are not entirely homogeneous among countries; however, strong similarities between clusters of European countries can be observed.

The results show that the approach is a useful and possibly complementary alternative to the existing analytical frameworks for assessing the impact of borrower-based macroprudential measures such as the LTV cap.

The major benefits of such an approach are the very few assumptions the method has to make on

the functional/distributional forms of the observed credit lending parameters, the lack of dependency on specific or complex data sources (any data source that enables the calibration of joint distributions would suffice), and the ability to easily incorporate additional features related to the behavioral response of all agents to such measures.

These results have clear macroprudential policy implications, since they allow gauging the potential impact of applying LTV caps both on the total number and the price impact of houses sold for each price segment of the house market. Such an analysis may also provide some guidance as regards banks' total credit contraction post-application of a specific measure that could be further used in informing conventional macroeconomic models.

1 Introduction

Housing is one of the most important sectors of modern economies, probably the largest asset class in the world, owing to its relationship with macroeconomic dynamics. The literature (see, for example, Reinhart and Rogoff (2013)) shows that real estate bubbles and bursts characterize almost all financial crises, with the recent episode of the Great Recession being no exception. However, because housing price characteristics such as illiquidity, locality, leverage or heterogeneity often render modelling demanding and difficult, researchers sometimes choose convenient shortcuts that represent a good approximation in most environments.

In this paper, we use an agent-based model to assess the effects of caps to the loan-to-value (LTV) ratio in the evolution of portfolio credit parameters, the impact on the provision of credit by banks and the evolution of housing prices following the application of such caps.

There is relatively extensive literature on housing markets. On the one hand, much conceptual work has started to appear related to macroprudential policy; some examples are Kuttner and Shim (2012), Nier *et al.* (2012), Kannan *et al.* (2012), and Christensen (2011). Mendicino (2012) shows that countercyclical LTV ratios in response to credit growth can smooth the credit cycle, whereas Unsal (2011) examines the relation between monetary policy and macroprudential regulation in an open economy DSGE model with nominal and real frictions. The author finds that macroprudential measures can usefully complement monetary policy. A useful recent paper summarizing the experiences with *ex ante* impact assessments of macroprudential instruments can be found in CGFS (2016).

On the other hand, research, though still scarce, is evolving on the empirical modelling side. Crowe *et al.* (2011) use state-level US data to find a positive relation between LTV at origination and subsequent property appreciation. Lim *et al.* (2011) evaluate the effectiveness of macroprudential instruments such as LTV caps in reducing systemic risk over time and across markets using data from 49 countries. Price (2014), as well as Bloor and McDonald (2013), use a Bayesian VAR to conduct *ex ante* counterfactual analyses prior to the introduction of borrower-based policies in New Zealand. Building upon the same approach, Cussen *et al.* (2015) conduct a micro simulation exercise based on loan-level data to quantify the impact of various caps on loan volumes in Ireland¹. Almeida *et al.* (2006) provide evidence that in countries with higher LTV ratios, house prices and demand for new borrowers are more sensitive to income shocks. Lamont and Stein (1999) find that in cities where a greater fraction of households have high LTV ratios, house prices respond more sensitively to economic shocks². For Korea, Igan and Kang (2011) find that LTV and debt-to-income ratio caps help to contain house price growth. Finally, Lambertini *et al.* (2011) highlight the importance of an expectations channel developing a model of the housing market that incorporates expectations-driven cycles, then showing that countercyclical LTV rules responding to credit growth can reduce the volatility of loans and the loan-to-GDP ratio.

¹Further related work for Ireland can be found in Hallissey *et al.* (2014), Lydon and McCarthy (2013) and Kelly (2011).

²A list of related studies includes, with no intent of being exhaustive, Gerlach and Peng (2005), Ahuja and Nabar (2011), Wong *et al.* (2011), Funke and Paetz (2012), and Wong *et al.* (2014).

The contribution of our paper, seen against this evolving strand of the literature, is to simulate a handful of simple models with agent-based techniques and to assess their efficiency in capturing the underlying dynamics. The intention is not to provide a fully-fledged analytical framework, rather a proof of concept on the suitability and efficiency of such behavior-capturing models to contribute to the impact assessment of macroprudential measures of this type. Therefore, the emphasis of the work is not on the numerical result itself but on the pros and cons of the modelling framework.

ABMs are not novel in the literature; some examples are Farmer (2014), Dawid et al. (2011), Colander et al. (2008), Gilbert et al. (2009) or LeBaron and Tesfatsion (2008). Our approach is based on Axtell et al. (2012 and 2014), who proposed a new and comprehensive model of the housing market of Washington, DC. This particular modelling approach is innovative in the literature in the sense that micro-level data is used to calibrate behavioral equations instead of postulating theoretical top-down behavioral rules. The main focus of their work was on demonstrating the causal relationship between leverage and the formation of a housing bubble.

Following the same overall approach in the use of multiple sources of micro data to elicit behaviors, Baptista et al. (2016) develop an ABM of the UK housing market to study the impact of macroprudential policies on key housing market indicators. This view enables them to tackle the heterogeneity in this market by modelling the individual behavior and interactions of first-time buyers, home owners, buy-to-let investors and renters from the bottom up, as well as to observe the resulting aggregate dynamics in property and credit markets.

In line with these works, in our paper the housing market is viewed as a universe of interacting heterogeneous agents comprising sellers, buyers and banks. Following autonomous decision rules, these agents interact directly with one another and with the economic environment, producing an overall economic outcome that emerges from complex interactions and that cannot be easily derived from the agents' objectives and behavioral rules.

In this respect, Table 1 compares our framework with the two aforementioned closest references. It is interesting to notice that our approach is the simplest one because we do not envisage to predict housing prices but to assess the impact of the application of borrower-based macroprudential measures. That is the reason why we do not consider neither investors nor renters in our model.

Since they largely influence the choice of simplifications in our approach, there are two specific characteristics of the task of measuring the impact of macroprudential measures that should be noted. First, the results are relative in the sense that answers are sought on metrics (credit provision, housing prices) with or without the application of caps. Second, for the time being, one-time-step models are more critical to design, since our focus is on the impact related to the application of the macroprudential measure and not on the convergence of more complex multi-hop ABM that tend to focus on forecasting the housing cycle.

	Baptista et al. (2016)	Axtell et al. (2014)	Buesa, Laliotis and Población (2019)
	<i>Macroprudential policy in an agent-based model of the UK housing market</i>	<i>An agent-based model of the housing market bubble in metropolitan Washington, DC</i>	<i>An Agent-based model for the assessment of LTV Caps</i>
Policy questions to address	(1) Booms and bust dynamics conditional on... ...size of rental/Buy-to-let sector ...different types of Buy-to-let investors (2) Qualitatively assess the effect of macroprudential policies, such as a loan-to-income ratio limit	Housing price dynamics (booms and busts)	Impact of borrower-based macroprudential measures (LTV, LTI caps)
Agent types	Households, banks, central bank	Households, banks	Households, banks and regulators (could be a central bank or a government)
Household types	First-time buyers, buy-to-let investors, renters, owner-occupier	Home buyers and sellers, investors	Home buyers and sellers
Role of bank agents	Supply of mortgage loans	Supply of mortgage loans	Supply of mortgage loans
Role of central bank agents	Set LTV and LTI	-	Set LTV caps
Role of government	-	-	Set LTV caps
Role of non-financial firms	-	-	-
Markets in the model	Housing and rental market	Housing and rental market	Housing market, mortgage loan market
Demographic features for households	Age (related: birth, death and inheritance)	-	-
Empirical calibration	Micro calibration: Household survey, housing market data; macroeconomic indicators	Micro calibration: Household survey, real estate transactions data, mortgage loans series	Any micro-data source that covers the required inputs (possibly HFCS database for European countries)

Note: Table 1 compares our framework with the two aforementioned closest references (Baptista et al. (2016) and Axtell et al. (2014)).

Table 1: Comparison between this paper and its closest peers in the literature.

The above characteristics are directed towards a “proof of concept” type of study, where one-time-step agent models are considered the basic building module, with the calibration design process focusing on the relative difference rather than absolute levels. As a result, we initially present simulated results based on stylized, yet pragmatic assumptions. Subsequently, we calibrate the probability distributions from empirical data. We use the second wave of the European Central Bank’s Household Finance and Consumption Survey, and we use a copula-based approach for the estimation of multivariate distributions of initial liquid wealth, debt-service-to-income ratios, LTV ratios and HMR Value at origination. Based on these real data probability distributions, we present

the distributions of properties that are actually traded in the auctions pre- and post-application of an absolute LTV ratio cap, as well as under the assumption of a proportionate cap³. Finally, we divide the sample by country in order to present the results at the country level.

The remainder of our paper is organized as follows: In Section 2, we sketch an outline of the model and its components. In Section 3, we present the results, first from a series of simulations and subsequently based on a model calibration on real survey data to demonstrate the impact of applying an LTV cap to empirical data. Section 4 concludes.

2 The agent-based model

ABMs are computational models in which heterogeneous agents interact directly with one another and with their economic environment, following autonomous decision rules.

Our approach is based on a very straightforward ABM where we model the interactions of sellers, buyers and banks within a computational framework that maps the economic environment (in the form of credit provision by banks and the applied LTV caps) into the emergence of a one-step housing market settlement/clearance state (specific properties are sold at specific prices during the one-step time interval). Therefore, the ABM can be viewed as a set of parallel-run property auctions, one for each seller, on the offered (by the sellers' universe) properties, where buyers' behavioral trends and banks' imposed financing constraints define the demand side. Since the problem is that of assessing the impact of imposing LTV caps, the whole computational model has to be set up as a relative/differential model in the sense that the results should be compared in the pre- and post- LTV cap cases.

In that context, we restrict ourselves to a single time step - how the housing market clears for one period with or without the cap - and we ignore any multiple period aspects - how housing prices evolve over time and how this is linked to the formation of price imbalances. This consideration is relatively scarce in the ABM literature, where more often the target is to predict the long-term equilibria for an entire housing cycle, that is, the multiple-period evolution of house prices and the respective imbalances.

2.1 The seller agents

The model assumes that there are N sellers at the beginning of the period, each of them offering a single property in the market. For each seller there is a parallel auction with all the buyers interested in buying the house. This implies that from a computational perspective, N parallel auctions would be needed in order to identify a final market clearing ratio, defined as the percentage of N that is

³Throughout this paper, it is assumed that an absolute LTV cap refers to a measure in accordance with which banks are not allowed to deviate from the imposed cap for any borrower. In contrast, a proportionate cap means that banks are allowed to exceed the imposed cap for a certain proportion of their newly originated exposures, i.e., some borrowers are indeed allowed to be granted loans in excess of the cap level, based on the banks' credit assessment.

finally sold in the market, and the “equilibrium” prices for those properties. The latter also entails the emergence of a settlement price for all properties: A transaction price for those sold and an average bid-ask price for those not sold.

The model is agnostic regarding the actual distribution of seller-asked prices. It only assumes that sellers uniformly cover the entire spectrum of the market, with no concentration on specific segments of buyers. More precisely, sellers can be ordered based on the asking price S_i , an ordering that is linked to the quality of the property, i.e. a higher price corresponds to a higher-quality property, with the term quality representing several quality-related features of the property, such as location, size, age, and proximity:

$$S = S_1, S_2, \dots, S_N, S_1 \leq S_2 \leq \dots \leq S_N$$

Calibration includes the setting of initial asking prices at levels that would ensure uniformity across the distribution of buyers, although alternative and more complex distributions of starting asking prices for the seller agents can be incorporated relatively easily.

For the sellers, the model also assumes a passive behavior in the sense that they start the auction with an asking price. If they are not “lifted”, that is, if the transaction does not originate by a corresponding buyer willing to pay the asking price, they will lower their asking price with probability p_d^S , trying to match a buyer agent in the auction until a limit of a fixed factor r percentage points from the initial asking price S_i is reached. In other words, there is a probability p_d^S that a seller, if not matched by a buyer at the initiation of the auction process, will gradually mark down its asking price from S_i to $S_i \times (1 - r)$ in an attempt to match a buyer: This is what we call an aggressive seller. Obviously p_d^S becomes important for calibration purposes, since it defines the supply side and the behavioral tendency of seller agents to mark down property prices.

2.2 The buyer agents

Buyer agents represent the households seeking to buy a property for their housing needs. A liquid wealth distribution is assumed for each buyer, a percentage of which may be used for the down payment of a mortgage loan. Each buyer is assigned an original LTV ratio which, combined with the maximum preferred down payment, results in the highest value of property up to which the agent can bid. Therefore, we assume a set of bids B from M buyer agents:

$$B = B_1, B_2, \dots, B_M, B_1 \leq B_2 \leq \dots \leq B_M.$$


where

$$B_k = \begin{cases} \min \left\{ \frac{D_k}{1-LTV_k}, a \times \omega_k \right\} & LTV_k < 1. \\ a \times \omega_k, & LTV_k \geq 1. \end{cases}$$

LTV_k represents the loan-to-value ratio associated with buyer agent k drawn from a probability distribution. Essentially, agents with an LTV ratio above 1 will bid as much as a multiple a of their liquid wealth ω_k . As for those having less than unity LTVs, they will choose between the latter amount and a multiple $1/(1 - LTV)$ of the down payment D_k , which is randomly drawn as a fraction of liquid wealth in the range (a_1, a_2) or calibrated on the basis of an empirical distribution.

The model also assumes that each buyer agent will always go for the highest quality property and bid in auctions for properties of higher value given the practical limitations imposed by its buying value price B_i . The model also assumes that with probability p_d^B , buyer k will also decide to participate in auctions where the starting asking price of the seller is smaller (up to a fixed factor g) than its original buying price B_k . This gives buyers the opportunity to react to excessive demand and competition by participating in the auctions of lower quality properties, always within the limits of a given distance g from the maximum quality they can attain with their buying price. Therefore, as in the case of p_d^S for seller agents, p_d^B for buyers represents their tendency to also go for lower quality due to overcrowding conditions. A high value of p_d^B reflects a higher density of buyer agents due to excessive demand.

The behavior of a buying agent during the initial phase, when no LTV limits/caps are considered, can be summarized by the decision tree in Figure 1.



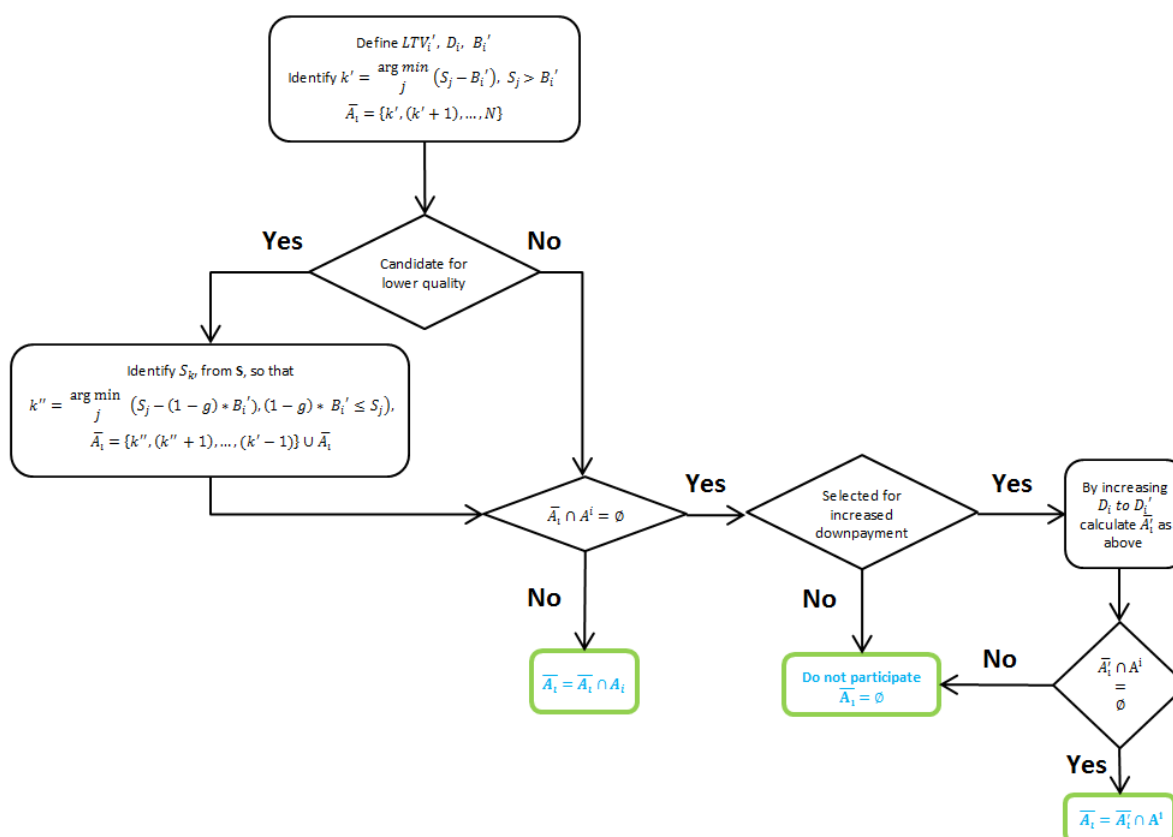
This figure shows with a decision tree the behavior of a buyer agent's during the initial phase, when no LTV limits/caps are considered.

Figure 1: Buyer Agent's Decision Tree for the No LTV cap case.

For the case of the agent's behavior following the application of LTV caps, an additional option is modelled: The choice, with fixed probability q_d^B , to raise the down payment for the mortgage so that the buyer ends up competing for properties that were within his quality reach before the implementation of the LTV cap (increased own participation). In other words, when the cap constrains the buyer agents to participate in auctions with properties below the range they had been allowed prior to the cap, they may opt with probability q_d^B to raise their down payment by increasing the part of household initial liquid wealth they consume in the purchase of the property. If such an increase does not allow them to participate in auctions for properties they would have participated in the no cap case, buyers remain inactive in the auctioning process.

Calibrating the value of q_d^B is a sensitive matter given that it is an artefact of our model through which buyers can liquefy their wealth. We assume that increasing the down payment is linked to a stable or improving financial situation; the European data used in the empirical exercise in Section 3.2 allows for some tentative calibration; more details can be found in Appendix 2.

With this addition, the buyer agent's behavior in the case of the application of LTV caps is summarized by the decision tree in Figure 2.



This figure shows with a decision tree the behavior of a buyer agent's in the case of the application of LTV caps.


Figure 2: Buyer Agent's Decision Tree for the LTV cap case.

2.3 The auction process

The N seller agents' prices $S = S_1, S_2, \dots, S_N, S_1 \leq S_2 \leq \dots \leq S_N$, are selected in a way that they are distributed across the entire spectrum of the buyer agents' bidding prices $B = B_1, B_2, \dots, B_M$. The easiest way to achieve this computationally is to divide the M buyers into N buckets, using N/M buyers for each bucket. By letting $B_i^*, i = 1, 2, \dots, N$ denote the maximum buying price within each bucket, the asking price of seller i is set as $S_i = B_i^* + \varepsilon$ where ε denotes a small positive constant.

Although such an assignment of asking prices is quite agnostic regarding the supply-side distributional characteristics, it corresponds to the relative strong assumption of a homogeneous market where sellers are distributed uniformly across the different quality segments. Based on the above description, Figure 3 schematically depicts the auctioning mechanism of N parallel auctions with the different types of sellers (aggressor and non-aggressor) and buyers (opting for lower-quality and non-opting ones).

The market clearing process can then be treated as a set of parallel auctions—one for each property—with the buyers' behavioral patterns fully defining the auctions' demand side. Computationally, this set of parallel auctions can be fully resolved by starting with the auction at the higher valued property and serially resolving lower-quality properties. This serialization enables the gradual clearance of both sellers and buyers, since buyers that have been successful in bidding a higher valued property can be removed from lower quality auctions.



This figure depicts the auctioning mechanism of N parallel auctions with the different types of sellers (aggressor and non-aggressor) and buyers (opting for lower-quality and non-opting ones).

Figure 3: Schematic representation of the parallel auctioning process.

The cap case is treated exactly the same as the additional behavioral characteristics for buyer agents accounting for flexibility, as discussed in the relevant section: Mutating to lower quality property

auctions and increasing down payments to reach the quality of the properties targeted by the agent in the no LTV cap case. Behavioral characteristics are maintained intact in the two runs: Aggressors are the same in both runs, and buying agents opting for the lower quality options will do the same for the LTV cap case.

Based on the above assumptions and behavioral models, the market clearing process is shown to converge for both cases (no cap and cap). It is important to point out that by controlling the number of buying agents relative to that of sellers, the probability that selling agents are willing to hit the bid of the best buyer if not lifted at the auction initiation and the probability that a buyer mutates to the auctions of lower quality can be used to calibrate the simulation process, provided that they are set to levels achieving the desired clearing ratio (percentage of properties sold before the LTV cap) and demand uplift indicator (percentage of properties sold at a price higher than the one asked by the seller): Essentially a combination of setting both the number of buying agents and their tendency to mutate.

It is relatively straightforward to calibrate control parameters so that the desirable clearing ratio and uplift ratio are attained in the no LTV cap case, and it is considered economically reasonable to restrain the initialization parameters to the two mentioned above. By way of the construction mechanism, the no LTV cap auction process results in an initial LTV distribution nearly identical to the desired one.

In this paper, we impose two different kinds of LTV caps: A simple absolute LTV cap and a more sophisticated one, the proportionate cap, in which banks are allowed to deviate from applying the cap only to a percentage of borrower exposures⁴. This second case requires the definition of a certain pecking order, based on which the deviation from the cap is applied to potential borrowers. In this case, we have assumed two different types of pecking order.

The first proportional cap benefits those with higher total wealth, which are the more likely to be granted a loan with an LTV exceeding what is allowed by the cap. To be more realistic, we added stochasticity to this selection process, allowing for some degree of randomness: A buyer is allowed to exceed the LTV cap with a probability that is analogous to his/her total wealth level; after selecting those that may be allowed to exceed the imposed cap, they are ranked based on their respective total wealth, and starting from the top ones, a serial selection process identifies the prospective buyers with the most total wealth until the exposure limit of the proportionate LTV cap is reached.

In the second type of proportionate cap, buyers closer to the median in total wealth will have a higher probability of receiving the loan with a higher LTV than what is allowed by the cap. Concretely, in this case, the probability of receiving the loan with a higher LTV than the one allowed by the cap is going to be inversely proportional to $\frac{|Wealth_{Total} - Wealth_{Median}|}{Wealth_{Median}}$, where $|\cdot|$ represents the absolute value. In both types of proportionate cap, this process may also involve some trial-and-error simulations in order to identify the number of borrowers that would be needed in order to reach the exposure percentage above the cap.

⁴Proportionate caps are common in OECD countries. See, for example, ESRB (2016) and Central Bank of Ireland (2015).

3 Results

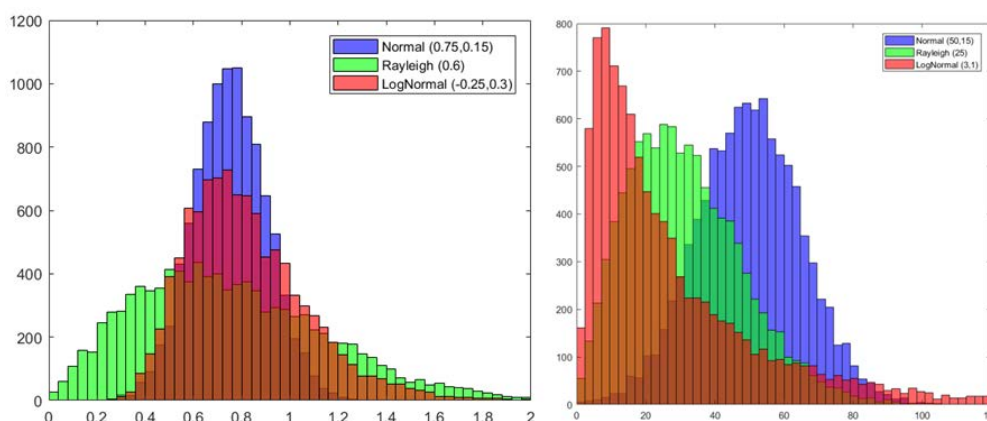
In this section, we elaborate on the results from our model. These results are important as a whole because they demonstrate that even simple ABMs that allow for the disaggregation of agents' basic behavioral characteristics may be used to assess or evaluate problems that can be viewed as problems of significant difficulty when demand and supply are treated in an average or aggregated manner.

A first exercise uses a collection of simulated probability distributions to study the main features of the responses and the sensitivities to selected parameter. In a second pass, we calibrate the distributions based on empirical European data extracted from the second wave of the HFCS survey combined with a copula methodology which helps us produce multivariate distributions of initial liquid wealth, total wealth, LTV ratio and property value at origination.

3.1 Simulated data

We carried out our study assuming that loan-to-value ratios and wealth follow three different types of probability distributions. The property value distribution function is derived from the LTV density.

Our starting points are a Gaussian with mean 0.75 and standard deviation 0.15 for LTV, and with mean 50 and standard error 15 for liquid wealth. Based on the Gaussians, we find densities from the Log-normal and Rayleigh⁵ families which cover the same range for both variables but exhibit different skewness and kurtosis. The chosen distributions are shown in Figure 4.



This Figure shows different assumptions for probability distribution of the LTV ratio (left) and the liquid wealth (right): Normal, Log-normal and Rayleigh. The brown area represents overlaps between the three distributions.

Figure 4: Probability distributions of LTV ratio (left) and liquid wealth (right).

The multiplicity of starting densities considered acts as a robustness exercise for our results but, more importantly, may be used to model groups of agents with different behavior in empirical exercises where data is scarce. For instance, one can think of a sample of home buyers obtained from a country


⁵The probability density function of a Rayleigh distribution with parameter ρ is $f(x|\rho) = \frac{x \exp\left(-\frac{x^2}{2\rho^2}\right)}{\rho^2}$.

where wealth is more uniformly distributed around the median and agents do not contract mortgage loans with high LTVs, even in the absence of a cap, owing to cultural aversion for overindebtedness.

The absolute LTV cap is set at 80% which corresponds to a reasonable level for most jurisdictions although, as will be discussed later (Table 2), by comparing the results for different LTV caps, useful insights into relative market impact can be extracted⁶.

The left hand side of Figure 5 summarizes the impact of the LTV cap on the distribution of 10000 buyer agents for the Gaussian case, which are allocated in buckets corresponding to the value of the property they are bidding for. The application of an LTV cap shifts the distribution of buyers' bids towards the lower end of the price range, since the cap becomes binding for a significant proportion of buyers if we also assume that there is no change in their household's liquid assets and their ability to come up with the required down payment. The inverse is true for higher-priced homes, where demand is relatively weak due to the cap application.

One interesting effect arises when combining differently shaped distributions for LTV and liquid wealth; on the right hand side of Figure 5 we exhibit the particular case of a Rayleigh density for LTV ratios and a Gaussian distribution for liquid wealth. As a consequence of the cap, a sizeable mass of buyers from the upper and lower end of the property range clusters more intensely around the average-valued homes, in a behavior mimicking that of a proportional cap penalizing buyers whose wealth is further from a centrality measure such as the median, which will be analysed later in this section.



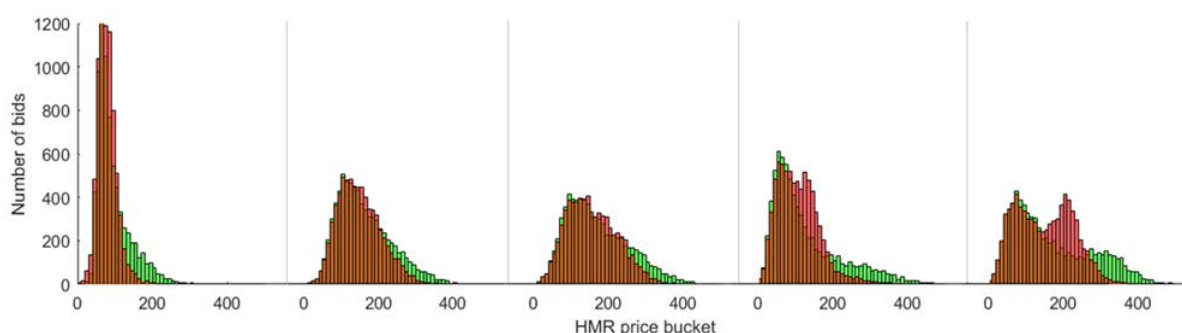
In this Figure we exhibit two particular cases. Case 1 (left) shows Gaussian density for LTV and liquid wealth whereas Case 2 (right) shows a Rayleigh density for LTV ratios and a Gaussian distribution for liquid wealth. The brown area represents overlaps between both distributions.

Figure 5: Pre- and post-cap distributions of buyer agents with simulated data.

From a statistical point of view, it is reasonable to reckon that the correlation between LTV and wealth distributions plays a major role in the shape of the results. Despite having agnostically assumed zero correlation in the previous exercise, most estimates for European countries from the

⁶Throughout this section and unless stated otherwise, we use the following calibration: $N = 7500$, $p_d^S = 0.2$, $p_d^B = 0.15$, $q_d^B = 0.3$, $r = 0.2$, $g = 0.1$, $a = 5$, $a_1 = 0.25$, $a_2 = 0.95$.


HFCS survey yield values of -0.3 on average; we therefore decide to compare the distribution of buyers across property buckets for five different correlations, which we show in Figure 6. Using the zero case as a benchmark, a greater positive correlation naturally shifts the distribution of buyers towards lower price ranges, since the cap constraint becomes binding for a significant proportion of households under the assumption that there is no change in their household's liquid assets and their ability to come up with the required down payment. In the most extreme case (right corner) the distribution becomes bimodal, as if the market were more uniform, still slightly polarized around two types of representative property: One high-quality for the wealthy households who sign mortgages with high LTVs, and one low-quality for poor households who become prudent and borrow as little as possible. For the negative correlation cases, the effect seems to be much more marked: The average traded price decreases abruptly while low-tier buyers with high-LTV loans flood the auctions for the cheapest properties.



In this Figure we compare the distribution of buyers across property buckets for five different correlations in the Gaussian-Gaussian case. Histograms correspond, in order, to correlation values of $\{-0.75, -0.3, 0.0, 0.3, 0.75\}$. The green area represents the no LTV cap case, the red area represents the 80% LTV cap case and the brown area represents the overlap of both distributions.

Figure 6: Distribution of buyers for different correlations between LTV and liquid wealth.

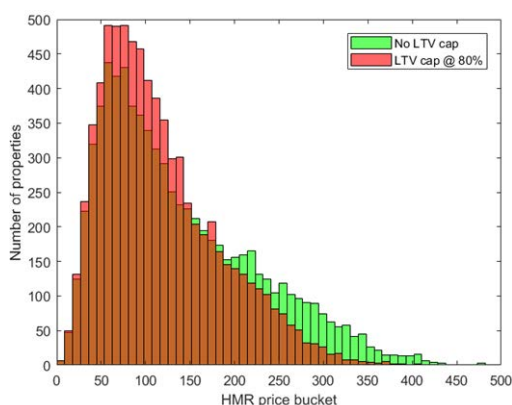
Figure 7 sketches the ordered (based on their respective bidding price) buyer agents and the bidding prices pre- and post- LTV caps. It also visualizes the non-uniform effect on prices that is also observed by comparing the distributions of Figure 5. This suggests that different segments of the housing price curve would be affected in a possibly different way. The magnitude of such effects would also depend on the households' initial liquid wealth distributions and the assumptions on their ability and willingness to increase down payments. Apart from the stylized approach used in the results presented here, where it is assumed that this behavior does not change post-application of LTV caps, the use of more-granular data to model possibly evolving behavioral patterns of buyers may result in significant changes in the way the demand side is comprised after the application of the cap.



Based on their respective bidding price, this figure sketches the ordered buyer agents and the bidding prices pre- and post- LTV caps.

Figure 7: Ordered buyer agents' bidding prices pre- and post-LTV cap.

Figure 8 presents the distributions of properties that were actually traded in the auctions, and it visualizes the potential impact of imposing an LTV cap on the number of properties sold and on the relative prices of these transactions. It is worth noting that the chart compares the pricing impact on individual properties that were traded under both market assumptions (LTV cap and no LTV cap), since the auctioning process may not warrant that the same individual properties changing hands after the application of caps. This figure is illustrative of the dampening effect of applying an LTV cap on both the amount and the prices that change hands on the basis of the specific agents' behavioral assumptions.

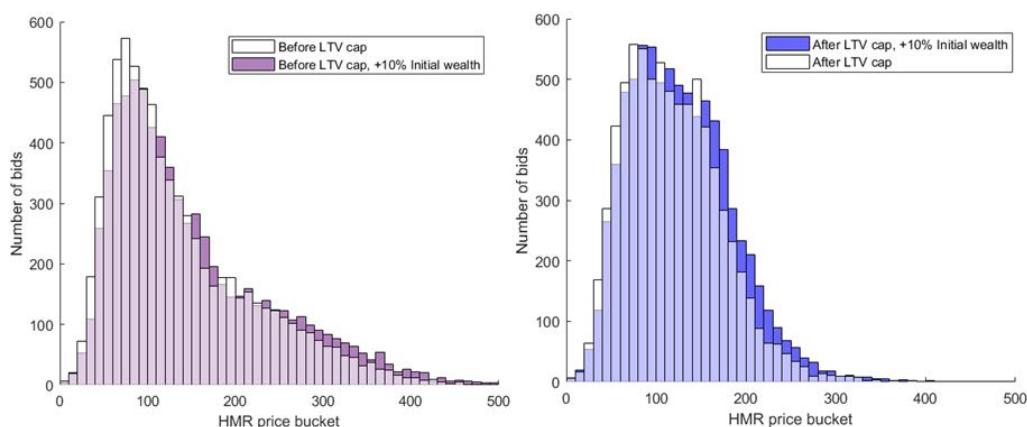


This figure presents the distributions of properties that were actually traded in the auctions for the Gaussian-Gaussian case using a correlation value of -30%. The brown area represents the overlap of both distributions.

Figure 8: Distribution of prices for sold properties pre- and post-LTV cap.

Taking into account the results from Figure 7 and 8, it is clear that the number of sold properties increases, especially in the price buckets around the mean value, and decreases in the higher-valued properties. One last consideration for the benchmark exercise of an absolute cap relates to how the distribution of agents pre- and post-cap may vary if liquid wealth increases by the same amount for

all buyers. Our ABM is not particularly sensitive with the benchmark calibration but points in a very intuitive direction: As shown in Figure 9, if wealth uniformly increases by one tenth of the initial value, the same buyers will be placed in auctions for mildly more expensive properties. Graphically, this entails a shift from the left to the right edges of the distribution.



This Figure shows the distributions of buyers for an overall 10% increase in wealth before (left) and after (right) the imposition of an LTV cap. In both cases there is an area which represents the overlapping areas for both distributions.

Figure 9: Effect of a global 10% increase in initial wealth on the distribution of buyers.

In addition to what has been presented above, a wider set of metrics and indicators can be collected and compared for the two cases in question. Table 2 summarizes the findings for simulated runs for different LTV cap values and market density parameters ($N/M, p_d^S, p_d^B$). Simulation results are presented for several runs that represent different levels of LTV caps (0.9, 0.85 and 0.8). The results in different clearance ratios in the first column of the table segments can be used to determine the combination of parameters that would be closer to the real conditions of the housing market. In other words, for each city or country, it is possible to obtain yearly data on sold properties and the stock of unsold properties to infer the clearance ratio and market density (N/M). With these real data, from Table 2, we can infer p_d^S, p_d^B , and other model initialization parameters.

The two segments of the table that refer to the cap and no cap cases also present statistics on the average traded price, the percentage of transactions that were secured by aggressive buyers, the percentage increase on the transaction price due to the aggressiveness of the buying agents and the split of the transaction price between credit provided by the bank and buyers' down payments.

Deriving metrics on the impact on housing prices is challenging due to the indexation process that might be required and the heterogeneity of the impact on different segments of the curve due to the "crude" way the LTV cap impacts the demand side of higher quality properties. On the other side, based on the assumptions made, estimating the variation of banks' credit provision levels and increased own participation levels is a straightforward process due to the detailed available data associated with the individual agents of the simulated model.

Nevertheless, in Table 2 we present two of those metrics; namely, the difference in credit provided by banks (capturing the credit supply impact) and the difference in property prices that were cleared in

The first noteworthy feature in Table 2 is that if we increase p_d^S and/or p_d^B , the clearance ratio increases because more houses are sold. It is interesting to note as well that, as expected, the higher the cap, the higher the percentage of buyers increasing down payments in order to buy a property. Conversely, if we increase p_d^B , the percentage of aggressors increases, but with a different intensity depending on the clearance ratio.

Differences in credit and price increases are more negative when there is a lower cap. This is unsurprising, since the lower the cap, the less credit is directed to buyers and, consequently, the less houses are sold. As explained above, the reduction in credit is much higher than the reduction in prices. Additionally, the difference in overall credit supply increases with p_d^S and the difference in prices increases with p_d^B , because when we increase p_d^S and/or p_d^B , there are more transactions, and consequently, the cap affects more transitions, which has a stronger impact on prices and credit.

As a last observation, we have to note that although the actual cap level appears to be the dominant factor in determining the impact on the demand, density parameters do have some importance in some cases. Therefore, we tend to believe that calibrating the model using real data would reflect prevailing market dynamics—especially concerning LTV distributions and the dominance of buyers or sellers—which may significantly contribute to the accuracy of the results.

The effect of an increase in r and g before the introduction of the cap is negligible; however, following the cap, the effect on the clearing ratio and the number of aggressor buyers is more visible, as Figure 11 illustrates; r is the markdown on prices the seller is willing to accept with probability p_d^s while g is the maximum distance that, with probability p_d^b , one buyer will deviate below the auction for the most expensive property he can afford. Both parameters are flexibility metrics; it is unsurprising that if they increase, the clearing ratio grows as there is more room for matching of buyers and sellers. The percentage of aggressor buyers, on the other hand, is only a function of the probabilities p_d .




Figure 10: Sensitivity of key auction variables to changes in r and g .

A similar logic applies when looking deeper into the mechanisms behind p_d^s and p_d^b in Figure 11: The more likely buyers or sellers are to haggle around the transaction price, the more properties will be sold thus increasing the clearing ratio. Besides, if p_d^s increases a buyer will be more likely to find himself bidding for a property which in principle he could not afford, so the percentage of aggressors

should fall because their desired effect is somehow coming from the supply side of the market. Finally, a larger p_d^b implies by definition more aggressor buyers, as seen in the first part of the red series on the left hand side; however, a non-linear effect arises after a threshold is reached, from where the number of aggressors falls, likely when the market is too saturated because of the rest of calibrated values.




Figure 11: Sensitivity of key auction variables to probabilities p_d^s and p_d^b .

3.2 Real data for European countries

In this sub-section, results based on the estimation of a multivariate distribution of initial liquid wealth, total wealth, LTV ratio and HMR Value at origination from a real data set⁸ are presented. To this end, we resort to the Household Finance and Consumption Survey (HFCS), a compilation of household surveys from European countries unified and collected by the European Central Bank.

Participating institutions, which are national central banks or national statistical institutes, conduct their own wealth surveys⁹. The HFCS provides the Eurosystem with harmonized micro-level data on euro area households' finances and consumption. The survey is conducted with a frequency of two to three years, with 2016 (second wave) as the most recent year that the current model calibration and assessment is based on.

The HFCS database is composed of questions that refer to the household as a whole. These questions are answered by one main respondent or by individual household members. Basic demographic information is requested in a personal questionnaire for all participating household members above sixteen years old. The survey part, covering household-level questions, encompasses real assets and

⁸Since there is enough data in the database for the four variables, we do not need to simulate property values from the LTV probability distribution function as we did in the simulation case.

⁹See https://www.ecb.europa.eu/pub/economic-research/research-networks/html/researcher_hfcn.en.html Link. A number of studies use the survey data for the primary purpose of measuring household vulnerability. See, for example, Gross and Población (2017), May et al. (2004), Johansson and Persson (2006), Vatne (2006), Herrala and Kauko (2007), Hollo and Papp (2007), Fuenzalida and Ruiz-Tagle (2009), Sugawara and Zalduendo (2011), Costa and Farinha (2012), Djoudad (2012), IMF (2012), Albacete and Lindner (2013), Albacete et al. (2014), Ampudia et al. (2016) and ECB (2014).

their financing, liabilities and credit constraints, private businesses and financial assets, intergenerational transfers and gifts and consumption/savings. Questions to individuals cover employment, future pension entitlements and labor-related income (other income sources are covered at the household level).

However, even though HFCS data are harmonized, since macroprudential policies are implemented differently in each country, HFCS empirical distributions might already be constrained by LTV, LTI and DSTI limits that each country has imposed¹⁰. This is a first limitation of our study because the framework would differ when imposing a tighter or looser LTV constraint in a country with existing caps. This issue cannot be solved because we cannot reconstruct the database assuming that there were no measures in place; fortunately, however, the active measures at the time of the survey were very few. We give more details in the country analysis section.

The second limitation of our empirical study comes from the fact that the HFCS dataset is a survey on outstanding, rather than new, loans. Hence, there may be a bias to cover longer-maturity (and therefore larger) mortgages. The impact of an LTV limit would instead be on new loans and thus would not affect the entire distribution.

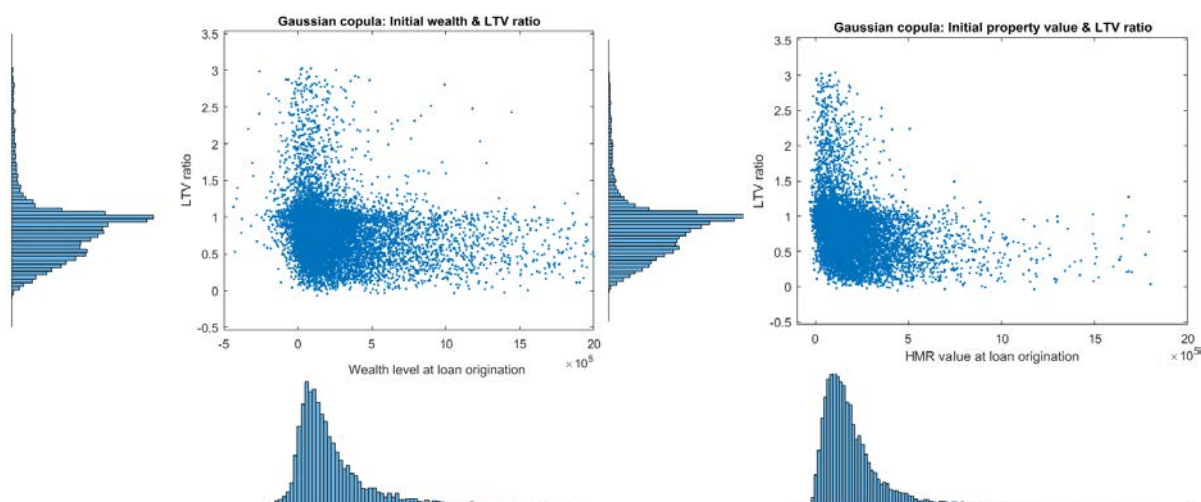
We use the well-known nonparametric approach (copula-based approach) for the estimation of multivariate distributions of initial liquid wealth, debt-service-to-income ratios, LTV ratios and HMR Value at origination. Our implementation of the nonparametric approach is via an accept-reject algorithm that can be sketched as follows. We start by estimating a four-variate Kernel distribution function (using an Epanechnikov Kernel) for each combination of. The four-variate probability density function $f(x)$ will be bounded by four pairs of minima and maxima. Uniform random numbers are strewn into these bounds, which delineate a four-dimensional polyhedron. Whenever a quartet of uniform random numbers falls under the joint probability density function, the quartet is accepted; otherwise, it is rejected. The resulting random numbers from the accepted draws will replicate the shape of the initial four-variate distribution¹¹. It is interesting to note that no distributional assumptions are imposed, neither on the marginal distributions nor on the copula that together constitute the joint distribution of the risk factors we examine.

In Figure 12, we present bivariate probability distributions based on the multivariate distributions of initial liquid wealth, debt-service-to-income ratios, LTV ratios and HMR Value at origination from the HFCS database.

Based on this multivariate probability distribution function, Figure 13 illustrates the distributions of properties actually traded in the auctions pre- and post-application of an LTV cap. We consider an absolute 85% limit.

¹⁰See, for example, ESRB (2016) to find some concrete examples.

¹¹For details about this general technique, which is an alternative to inverse CDF transform-based methods, see, e.g., Ross (1990). The alternative to the so-called smooth bootstrap version we describe here is what one could call 'plain' bootstrap, which would not involve a univariate or multivariate Kernel estimator in the first step but would resample directly from historical data. The reason for considering a smooth bootstrap is to avoid replicating possibly fine though spurious details in the data, which is a concern in short samples in particular.



Using the HFCS database we present the bivariate probability distributions based on the multivariate distributions of initial liquid wealth, debt-service-to-income ratios, LTV ratios and HMR value at origination.

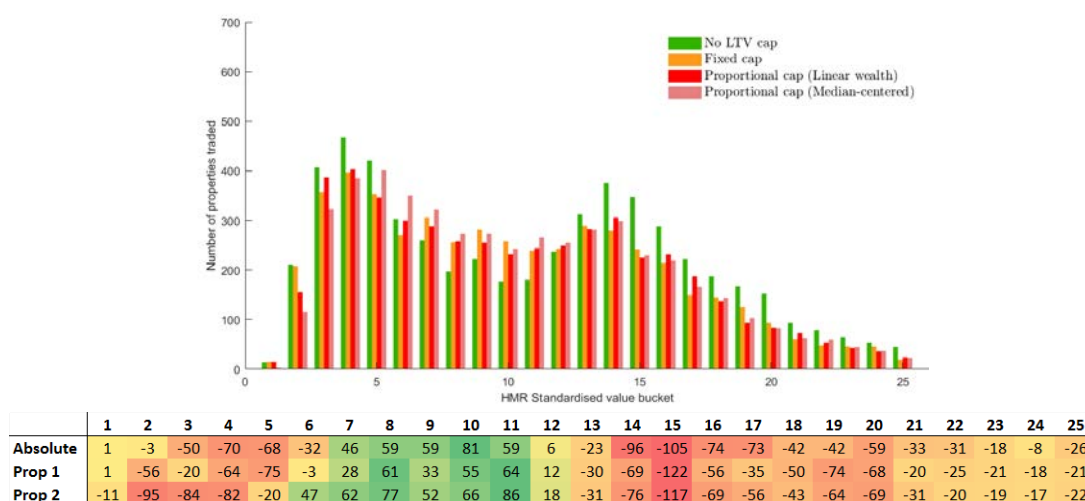
Figure 12: Probability distributions for aggregate HFCS data.

On one hand, in the previous section we have explained that in both proportionate caps the probability of exceeding the LTV cap is based on the total wealth. In the first one, the higher the wealth, the higher the probability whereas in the second one this probability is higher in medium wealth households. On the other hand, from Figure 12 we observe that on aggregate high wealth is related to moderate LTV values (seldom above 100%) whereas mid- and low-level wealth are linked with high LTVs. From the standpoint of our exercise, buyers with their LTV ratio in the medium range (and the upper part of the income distribution) are the most likely to exceed the cap in the wealth-related proportionate case; such effect is likely to be present under a median-centered cap, too, but now buyers to the left of the median (not very wealthy, still with moderately low LTV) can also beat the cap.

For the absolute cap, most of the market movement is clustered in the low-quality housing range, where buyers with high LTV will purchase fewer properties. This evidence can hint the existence of subprime layers of loans in the European mortgage market, which is something with quite substantial macroprudential policy implications. In parallel, wealthier agents with high LTVs will shift to cheaper segments. All in all, the combination of both yields a displacement of the transaction mass to the left of the distribution. However, as we will explain in the next section, these results are not homogeneous among countries.

The application of an LTV cap may have different impacts in different segments of the housing market, which, can lead to undesirable consequences. That is why it is reasonable to consider the possibility of a proportionate cap instead, as explained in the second section. As stated in Section 2, the proportionate cap requires the definition of a certain pecking order based on which the cap is applied to clients. In the first case, we have assumed that banks show a preference to wealthier buyers. By doing so, banks comply with the intuitive creditworthiness principle, which suggests that potential borrowers with higher total wealth are more likely to be granted a loan with an LTV that exceeds the cap threshold. In the second variant, buyers closer to the median in wealth will have

a higher probability of receiving the loan with a higher LTV than the one allowed by the cap. This illustrates a more expansive policy for banks that are trying to distribute the excessive leverage to a larger number of potential borrowers of lower average wealth.



This figure illustrates the distributions of properties actually traded in the auctions pre- and post-application of three LTV caps (one absolute and two proportionate caps). The table under the graph reflects the absolute change in traded properties with respect to the no-cap case.

Figure 13: Distribution of prices for sold properties pre- and post-LTV cap - HFCS data.

The aggregate HFCS results for the case of proportional caps are misleading, in the sense that they suggest that all three caps generate almost identical responses. For this reason, we postpone the analysis for the two cases to the following subsection, where much clearer conclusions can be drawn.

3.2.1 Country analysis

The HFCS database comprises twenty countries. However, as can be seen in Table 3, we do not have the same number of households with all the required information for each country. Moreover, in some of them (FI, HU, LV, MT, SI), the number of data-sufficient households is too small to carry out a proper analysis, and consequently, they had to be excluded¹².

For the remaining countries, we compared the period of data collection for each national survey with applications of LTV caps by the relevant authority. Cyprus, Latvia and Poland introduced limits as the survey was taking place; only the Netherlands, for which data was gathered in January and February 2014, had a loan-to-value cap of 100% since 2012¹³.

¹²In the case of Finland, no households at all with all the required information are available.

¹³More precisely, the Netherlands introduced a cap of 106% in 2012 which progressively lowered 100% in 2018. Among the 589 valid households in the survey, only 4% contracted loans between 2012 and 2014.

Country code	# Of valid households	Country code	# Of valid households
AT	372	IE	1566
BE	508	IT	546
CY	362	LU	507
DE	827	<i>LV</i>	<i>33</i>
EE	307	<i>MT</i>	<i>114</i>
ES	1016	NL	589
<i>FI</i>	<i>0</i>	PL	262
FR	2260	PT	1756
<i>GR</i>	<i>198</i>	<i>SI</i>	<i>139</i>
HU	41	SK	192
		Total	11595

This table presents the number of households in HFCS database with all information available by country. Rows in red represent countries ruled out of the analysis due to an insufficient number of households.

Table 3: Number of households in HFCS database with all information available by country.

Before showing what our agent-based model has to say on the application of caps across countries, a closer look at the initial loan-to-value ratio distributions¹⁴ for each geography, shown in Figure 14, reveals notable differences in a number of dimensions; In terms of the third moments, densities are clearly left-skewed for Austria, Cyprus, Germany, and right-skewed for Spain, Ireland and Portugal; regarding kurtosis there are some cases of platykurtic distributions, most noticeably Poland and Slovakia, in contrast with Spain and Portugal which are more leptokurtic. These features, along with those of wealth distributions¹⁵, will shape and cluster the results considerably.

In Figure 15, we present the country distributions of prices for traded properties pre- and post-application of all three versions of an LTV cap, grouped in 25 standardized value buckets for better comparison. At the country-aggregate level, the total number of houses sold decreases following the introduction of whichever cap; Besides, the proportional caps have a distinctively stronger effect than the fixed alternative in the lowest price buckets: Buyers under the wealth-linear cap will have virtually zero probability of bypassing it because of their low level of wealth; if under the median-centered alternative, they can do no better as they are too far from the median value.

Our model, on the contrary, is moderate in the impacts for the highest-priced property buckets. Some wealthy buyers may, of course, have contracted a loan with a high LTV ratio, be it because of cultural reasons or for having been able to jump over the cap. Nevertheless, the absolute decline in the traded properties for these market segments is very low or negligible with respect to other quantiles of the distribution.

¹⁴A comparison of empirical and copula-generated densities can be found in the Appendix.

¹⁵We also conducted a preliminary analysis on the cross-country wealth distributions, but the differing supports and the similarity of shapes do not justify including it in the final conclusions.


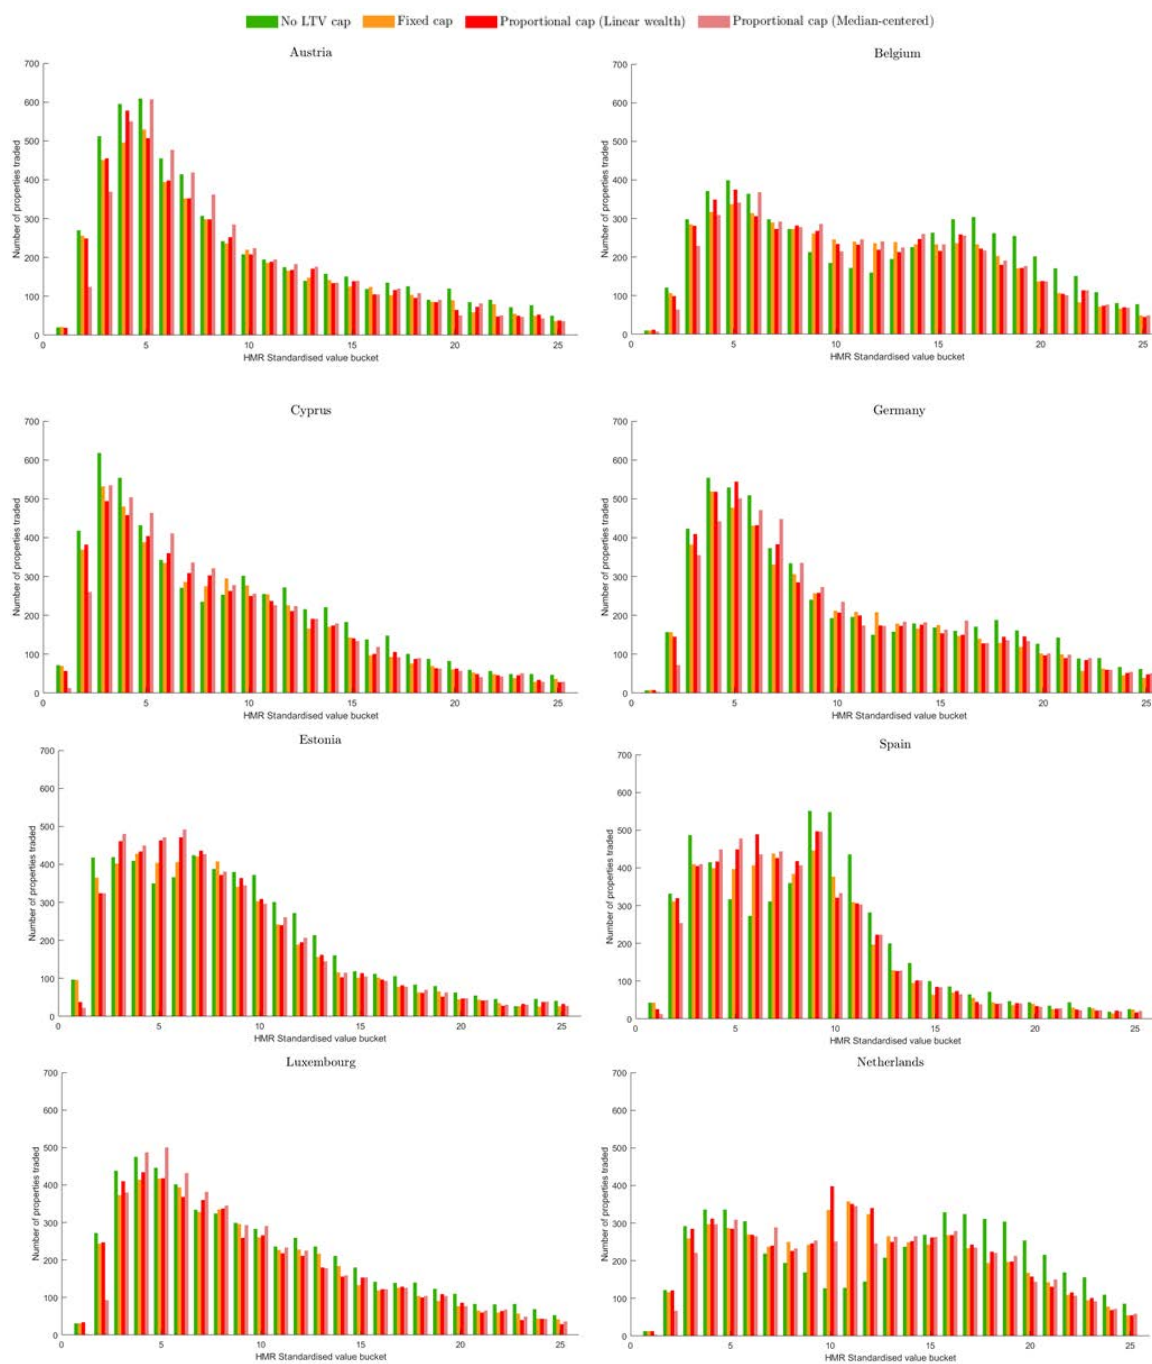


Figure 14: Generated loan-to-value ratio distributions by country.


The last stylized fact of the country-level responses to the cap is what was already noticed at the aggregate HFCS level: There is a tendency for wealthy buyers with not-so-low LTV ratios, possibly close to the cap value, to shift towards cheaper properties following the introduction of the measure; this behavior results in a large mass shifting to the left of the distribution, hence lowering the average trading price.



In this Figure we present the country distributions of prices for traded properties pre- and post-application of all three versions of an LTV cap (one absolute and two proportionate caps).

Figure 15: Country distributions of prices for sold properties pre- and post-LTV caps.

(Figure 15 continued)



While reading succinctly through the LTV ratio distributions, we reckoned that the variety of their shapes had to be relevant for the results. Indeed, comparing the series of absolute declines in traded properties by cap, country and bucket has allowed us to distinguish five groups of countries:

1. *“Price-symmetric” (BE,IE,NL)*: The effect of all three caps is similar in magnitude; houses traded in the cheapest and most expensive property segments decrease by a similar, considerable amount, while they increase symmetrically in the central buckets. All three countries have considerably more households above the 80% and around the 100% LTV ratio, including wealthy agents. In fact, the first quartile of the wealth distribution in NL and BE is almost empty, suggesting strong appetite for saving.

The result is a net increase in medium segments coming from a net decrease in the extreme ones.

2. *“Wannabe-targeting” (EE,ES,PT)*: Again, the effects of the three variants of the cap do not differ a lot, but now lower-priced buckets suffer the cap much less and the movement is concentrated in the medium ones: Within them, the cheapest are traded much more coming from the highest-

valued ones being traded much less. An interpretation is that wealthy households in the middle of the distribution have an appetite for high-LTV loans¹⁶. Meanwhile, the richest buyers borrow at low LTVs.

The result is a negligible change in medium segments due to offsetting positive/negative within variation.


3. “Subprime-targeting” (AT,FR,IT,LU): The absolute and wealth-linear proportionate cap behave similarly while the median-centered one is different and has a stronger effect. Only the cheapest property buckets are affected. All in all, this suggests a flat distribution of LTVs combined with a left-skewed wealth distribution, implying a large number of households with little wealth and high LTVs.

The result is a net decrease in low segments.

4. “Flat” (CY,PL,SK): Similar to case 3, but now the absolute cap has a smaller effect than the two proportional ones, which are very close to one another. Such feature points to a left-skewed, leptokurtic wealth distribution: The effect on the low-tier buyers in the case of the proportionate caps is very similar because the median of the distribution is centered in the major part of the mass.

5. “Hybrid” (DE,GR): These two countries exhibit features from the other four categories. For illustration, Greece is somehow “price-symmetric” in the sense that effects on cheapest and best properties are symmetric in magnitude, but also “subprime-targeting” owing to the more pronounced effect of the median-centered cap in the cheapest buckets. Note that, for example, the wealth distribution for Greece is virtually symmetric while the LTV ratio density is quite flat.

An example of each type of country is depicted in Figure 16. The full results can be found in the Appendix, along with a comparison of the loan-to-value ratio distributions before and after the application of a proportional cap. Intuitively, the wealth-proportional version fattens the right tail of the density while the median-centered one renders the latter more symmetric.



In this figure we present an example of each type of country by cap type described above: Absolute (orange), Wealth-linear proportionate (Dark red), Median-center proportionate (light red). Results are presented in differences with respect to the no cap case.

Figure 16: Absolute variation in properties traded by bucket for all types of countries.

¹⁶A pattern that was, indeed, very frequent among Spanish households preceding the burst of the real estate bubble in 2008.

The results we have discussed in this section lead us to think that our ABM is a useful tool to assess the impact of macroprudential policy measures at the country level, and that it can be a robust complement to the existing, more parsimonious analytical frameworks. Moreover, our findings stress the need for careful implementation of policies, regardless of their simplicity, the effects of which can vary considerably across agents as a function of their endowments (say wealth), appetite for credit (say LTV) or flexibility of their preferences (say willingness to accept purchasing lower-quality housing).

4 Conclusions

In this paper, we propose a simple model relying upon agent-based techniques to assess the impact of a cap to the loan-to-value ratio.

Results based on simulated data are presented first, assuming that initial liquid wealth, total wealth, LTV at origination and property value parameters following three types of probability distributions in which initial liquid wealth and LTV at origination can be positively and negatively correlated. The application of an LTV cap naturally shifts the distribution of buyers towards lower price ranges, since the cap constraint becomes binding for a significant proportion of households under the assumption that there is no change in their household's liquid assets and their ability to come up with the required down payment. The inverse is true for higher-priced properties, where demand is relatively weak due to the cap.

After this simulation exercise, the relevant probability distributions are calibrated on actual European data. In that context, the second wave of the Household Finance and Consumption Survey (HFCS) is used. We also deploy a copula-based approach for the estimation of multivariate distributions of initial liquid wealth, debt-service-to-income ratios, LTV ratios and HMR Value at origination. High-value properties are not necessarily linked with high LTVs in the data; in fact, high LTVs correspond mainly to the low and medium value of the house price range. Hence, when we impose a cap on LTV, low and medium value houses are more affected than high value houses. This evidence can also be considered as revealing the presence of “subprime” mortgage loan segments in the European loan markets, a potentially significant feature with macroprudential policy implications.

When LTV cap impacts are calculated country by country, considerable heterogeneity arises. We are able to distinguish between five groups of countries owing to their reaction to the three versions of the cap and the impact on the different price segments.

As stated above, based on the results, we think that the approach is a useful and complementary alternative to the existing analytical framework for assessing the impact of macroprudential borrower-based measures such as LTV caps. The major benefits are the very few assumptions our method has to make on the functional/distributional forms of observed credit lending parameters and its ability to incorporate, even in a probably unsophisticated fashion, features related to the behavioral response of borrowers to such measures. Moreover, due to the simplicity of the model, many simple extensions can be added. For example, sharper sequential mechanisms for the property auctions or more than one

time step. The vast amount of empirical data available may also allow for country-level calibration.

All in all, the results presented in this study have clear macroprudential policy implications, since they allow us to gauge the potential impact of applying LTV limits to both the number and the price of houses sold for each segment. Such an analysis may also provide some guidance as regards the banks' credit contraction post-application of a specific cap.

References

- [1] Ahuja, A. and Nabar, M. Safeguarding banks and containing property booms: Cross-country evidence on macroprudential policies and lessons from Hong Kong SAR. IMF Working Paper No. 11/284, December 2011.
- [2] Albacete, N., Lindner, P. Household vulnerability in Austria – A microeconomic analysis based on the Household Finance and Consumption Survey. OeNB Financial Stability Report 25. pp. 57–73, June 2013.
- [3] Albacete, N., Eidenberger, J., Krenn, G., Lindner, P., Sigmund, M. Riskbearing capacity of households – Linking micro-level data to the macroprudential toolkit. OeNB Financial Stability Report 27. pp. 95–110, June 2014.
- [4] Almeida, H., Campello, M., and Liu, C. The financial accelerator: Evidence from international housing markets. *Review of Finance*, 10(3), 1-32, 2006.
- [5] Ampudia, M., van Vlokhoven, H. and Żochowski, D. Financial fragility of euro area households, *Journal of Financial Stability*, Vol. 27, 250–262, 2016.
- [6] Axtell, R., Conlee, B., Farmer, J. D., Geanakoplos, J., Howitt, P., Goldstein, J., Hendrey, M., Palmer, N. M., and Yang, C.Y. Getting at Systemic Risk via an Agent-Based Model of the Housing Market. *American Economic Review* 102(3), 53–58, 2012.
- [7] Axtell, R., Carrella E., Conlee, B., Farmer, D., Geanakoplos, J., Jon Goldstein, J., Hendry, M., Howitt, P., Kalikman, P., Masad, D., Palmer, N. and Yang, C. An agent-based model of the housing market bubble in metropolitan Washington, DC. 2014.
- [8] Baptista, R., Farmer, J.D., Hinterschweiger, M., Low, K., Tang, D., and Uluc, A. Macroprudential policy in an agent-based model of the UK housing market. Staff Working Paper No. 619, 2016.
- [9] Bloor, C. and McDonald, C. Estimating the impacts of restrictions on high LTR lending. Reserve Bank of New Zealand Analytical Note No. AN2013/05, October 2013.
- [10] Central Bank of Ireland. Information Note: Restrictions on residential mortgage lending. Central Bank of Ireland, 2015.
- [11] CGFS. Experiences with the ex-ante appraisal of macroprudential instruments. CGFS Paper No. 56, July 2016.
- [12] Christensen, I. Mortgage debt and procyclicality in the housing market. Bank of Canada Review, 2011.
- [13] Colander, D., Howitt P., Kirman A., Leijonhufvud A., and Mehrling P. Complexity and Dynamics in Macroeconomics: Alternatives to the DSGE Models - Beyond DSGE Models: Toward an Empirically Based Macroeconomics. *American Economic Review: Papers & Proceedings*, 2008.
- [14] Costa, S., Farinha, L. Households' indebtedness: a microeconomic analysis based on the results of the households' financial and consumption survey. Banco de Portugal Financial Stability Report, pp. 133–157, May 2012.

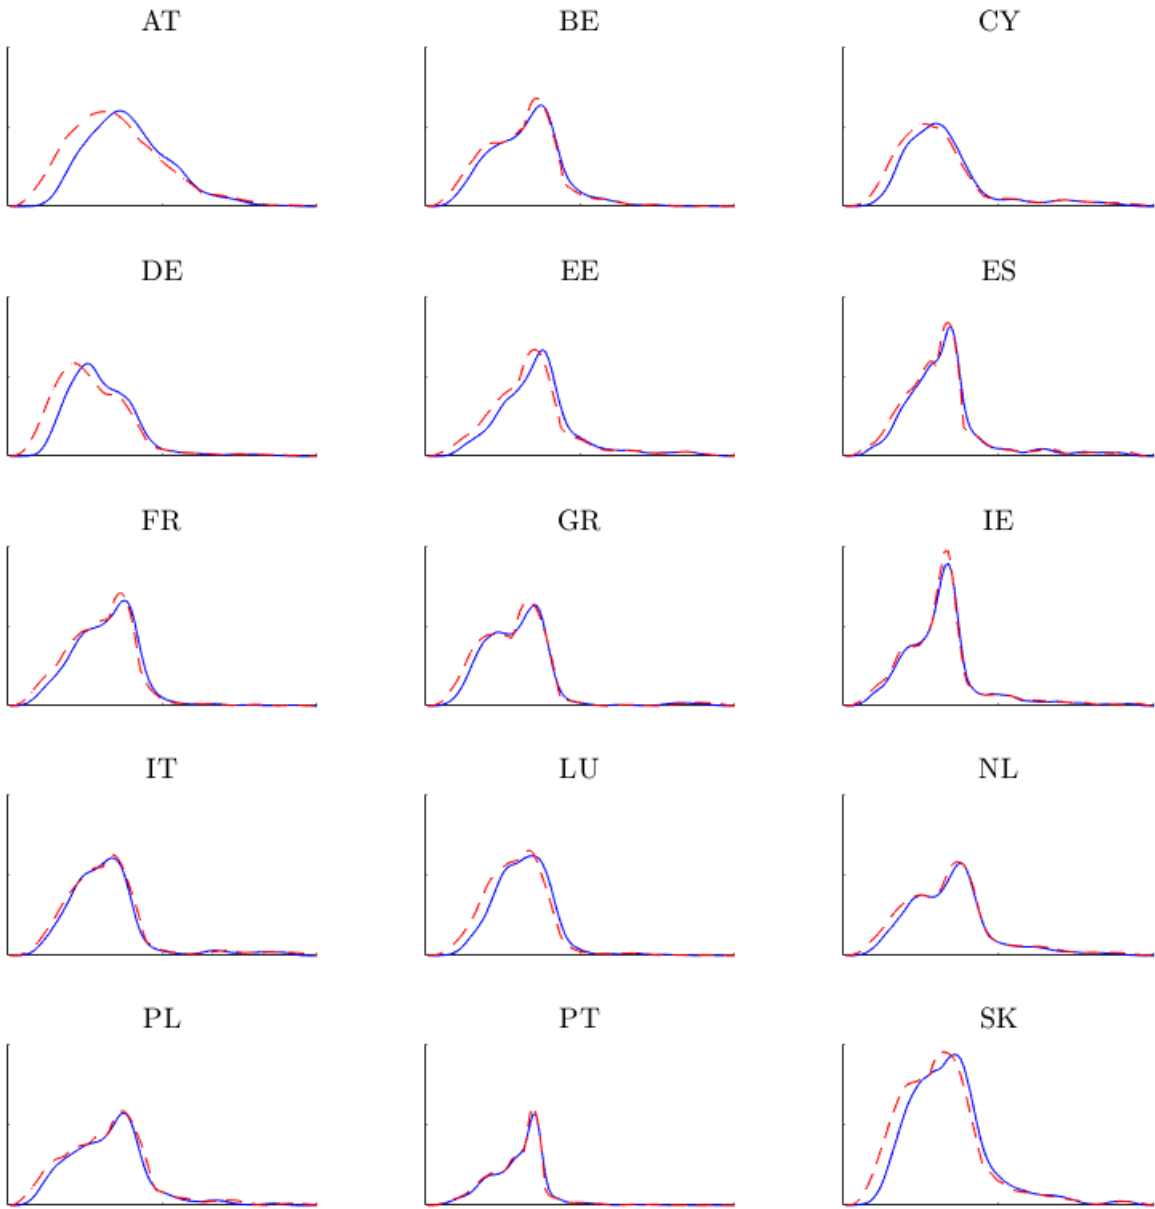
- [15] Crowe, C., Dell’Ariccia, G., Igan, D., and Rabanal, P. Policies for macro financial stability: Options to deal with real estate booms. IMF Staff Discussion Note 11/02, February 2011.
- [16] Cussen, M., O’Brien, M., Onorante, L., and O’Reilly, G. Assessing the impact of macroprudential measures. Central Bank of Ireland Economic Letter Series, 2015(3), 2015.
- [17] Dawid, H., Gemkow, S., Harting, P., van der Hoog, S., and Neugart, M. The Eurace@Unibi Model: An Agent-Based Macroeconomic Model for Economic Policy Analysis. October 2012.
- [18] Djoudad, R. A framework to assess vulnerabilities arising from household indebtedness using microdata. Bank of Canada Discussion Paper 2012–3, February 2012.
- [19] ECB. The financial vulnerability of euro area households – Evidence from the Eurosystem’s Household Finance and Consumption Survey. ECB Monthly Bulletin, pp. 79–90, November 2014.
- [20] ESRB. A Review of Macroprudential Policy in the EU in 2015. ESRB, May 2016.
- [21] Farmer, D. Agent Based Modelling in Economics: promises and challenges. European Central Bank, December 2014.
- [22] Fuenzalida, M.C., Ruiz-Tagle, J.V. Household financial vulnerability. Central Bank of Chile Working Paper No. 540, December 2009.
- [23] Funke, M. and Paetz, M. A DSGE-based assessment of nonlinear loan-to-value policies: Evidence from Hong Kong. Bank of Finland BOFIT Discussion Papers 11/2012., April 2012.
- [24] Gerlach, S. and Peng, W. Bank lending and property prices in Hong Kong. *Journal of Banking and Finance*, 29(2), 461:481, February 2005.
- [25] Gilbert, N., Hawksworth J., and Swinney, P. An Agent-Based Model of the English Housing Market. Association for the Advancement of Artificial Intelligence, 2009.
- [26] Gross, M. and Población, J. Assessing the efficacy of borrower-based macroprudential policy using an integrated micro-macro model for European households, *Economic Modelling*, 2017.
- [27] Hallissey, N., Kelly, R., and O’Malley, T. Macroprudential tools and credit risk of property lending at Irish banks. Central Bank of Ireland Economic Letter Series No. 10/EL/14, 2014(10), 2014.
- [28] Herrala, R., Kauko, K. Household loan loss risk in Finland – Estimations and simulations with micro data. Bank of Finland Research Discussion Paper No. 5/2007.
- [29] Hollo, D., Papp, M., December 2007. Assessing household credit risk: Evidence from a household survey. Magyar Nemzeti Bank Occasional Papers No. 70, 2007.
- [30] Igan, D. and Kang, H. Do loan-to-value and debt-to-income limits work? Evidence from Korea. IMF Working Paper No. 11/297, December 2011.
- [31] IMF. Spain: Vulnerabilities of private sector balance sheets and risks to the financial sector. Technical Note in IMF Country Report No. 12/140, June 2012.
- [32] Johansson, M.W., Persson, M. Swedish households’ indebtedness and ability to pay – a household level study. Sveriges Riskbank Econ. Rev. 3 (December), 24–41, 2006.

- [33] Kannan, P., Rabanal, P., and Scott, A.M. Monetary and macroprudential policy rules in a model with house price booms. *The B.E. Journal of Macroeconomics*, 12(1), 2012.
- [34] Kelly, R. The good, the bad and the impaired | A credit risk model of the Irish mortgage market. Central Bank of Ireland Research Technical Paper, No 13/RT/11, 2011.
- [35] Kuttner, K. and Shim, I. Taming the real estate beast: The effects of monetary and macroprudential policies on housing prices and credit. Property Markets and Financial Stability, Conference Volume of Reserve Bank of Australia Annual Conference, 2012.
- [36] Lambertini, L., Mendicino, C., and Punzi, M. Leaning against boom-bust cycles in credit and housing prices. Bank of Portugal Working Paper No. 8-2011, March 2011.
- [37] Lamont, O. and Stein, J. Leverage and housing price dynamics in US cities. *RAND Journal of Economics*, 30(3), 498:514, 1999.
- [38] LeBaron, B., and L'Tesfatsion, L. Modelling Macroeconomies as Open-Ended Dynamic Systems of Interacting Agents", *American Economic Review: Papers & Proceedings*, 98(2), 2008.
- [39] Lim, C., Columba, F., Costa, A., Kongsamut, P., Otani, A., Saiyid, M., Wezel, T., and Wu, X. Macroprudential policy: What instruments and how to use them? Lessons from country experiences. IMF Working Paper No. 11/238, October 2011.
- [40] Lydon, R. and McCarthy, Y. What lies beneath? Understanding recent trends in Irish mortgage arrears. *The Economic and Social Review*, 44(1), 117-150, 2013.
- [41] May, O., Tedula, M., Young, G. British household indebtedness and financial stress: a household-level picture. Artic. Bank Engl. Q. Bull, 2004.
- [42] Mendicino, C. Collateral requirements: Macroeconomic fluctuations and macro-prudential policy. Banco de Portugal Working Paper 11/2012, April 2012.
- [43] Nier, E., Kang, H., Mancini, T., Hesse, H., Columba, F., Tchaidze, R., and Vandebussche, J. The interaction of monetary and macroprudential policies: Background paper. IMF Staff Paper, December 2012.
- [44] Price, G. How has the LVR restriction affected the housing market: A counterfactual analysis. Reserve Bank of New Zealand Analytical Note No. AN2014/03, May 2014.
- [45] Reinhart, C. and Rogoff, S. Banking crises: An equal opportunity menace. *Journal of Banking & Finance*, 37, 4557–4573, 2013.
- [46] Sugawara, N., Zalduendo, J. Stress testing Croatian households with debt: implications for financial stability. World Bank Policy Research Working Paper No. 5906, December 2011.
- [47] Tesfatsion, L. S. and Judd, K. L. Handbook of computational economics. North Holland 2006.
- [48] Unsal, D.F. Capital flows and financial stability: Monetary policy and macroprudential responses. IMF Working Paper No. 11/189, August 2011.
- [49] Vatne, B.H. How large are the financial margins of Norwegian households? An analysis of micro data for the period 1987–2004. Artic. Norges Bank Econ. Bull. 77 (4), 173–180, 2006.

- [50] Wong, E., Fong, T., Li, K., and Choi, H. Loan-to-value ratio as a macroprudential tool: Hong Kong's experience and cross-country evidence. Hong Kong Monetary Authority Working Paper No. 01/2011, February 2011.
- [51] Wong, E., Tsang, A., and Kong, S. How does LTV policy strengthen banks' resilience to property price shocks: Evidence from Hong Kong. Hong Kong Monetary Authority Working Paper No. 03/2014, February 2014.

Appendix 1: Accuracy of copula-based LTV distributions *vis-à-vis* HFCS data

In order to provide a metric how close the simulated cross-country LTV distributions are to the observed one in the data, we have created a matrix of plots with the kernel density estimators for both. As can be seen, the simulation mechanism succeeds in replicating the densities in the original data.



In this Figure we provide a set of plots with the kernel density estimators for both the simulated and the HFCS data (dashed red for the HFCS, blue for the simulated).

Appendix 2: Details on the HFCS data

A2.1. Data cleaning and preparation

In a first step (say “quality assurance”), we remove from the sample every observation for which any of the following is true:

- DL2100i (Has mortgage payments) is not 1.
- DA1110i (Has HMR) is not 1.
- fHB0800 (Flag for property value at acquisition) is “not imputed”, “originally not collected” or “originally no answer”.
- fHB1401 (Flag for initial amount borrowed) is “not imputed”, “originally not collected” or “originally no answer”.
- DN3001 (Net wealth) is blank.
- HB0800 (Property value at acquisition) is blank.
- DL2110 (Mortgage payments for HMR, flow) is blank.
- DA1110 (Value of HMR) is blank.

Once the data is free from blanks and missing values, we filter the observations that:

- Have an LTV ratio over 300%.
- Have a debt service-to-income ratio above 50% or below 0%.
- Have a net wealth-to-HMR value ratio of above 50% in absolute value.

After the data cleaning process, we are left with 11595 observations.

A2.2. Calibration of q_d^B

We have tentatively calibrated this parameter using some aggregate information from the HFCS. q_d^B represents the probability of a buyer increasing his down payment; we deem reasonable to assume that this will only happen if its financial situation is sufficiently stable or, moreover, likely to improve. For this purpose, we use four indicator variables:

- HNB1700 (“Household makes extra mortgage payments over contractual amount”).
- HNK0400 (“Household expects the overall economic situation to improve”).
- HNI0700 (“Household expects to have more savings next year”).
- (1-PNE2800x) (“Household expects the work situation *not* to worsen in the near future”).

We calculate the probabilities of positive responses conditional on data availability for all 4 variables, then set q_d^B equal to their average. We obtain a value of 0.3024, which is what we use in the paper.

Appendix 3. Detailed country results

The following tables show the absolute change in traded properties for all three cap cases by standardised HMR value bucket.

Absolute	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
AT	1	-14	-61	-99	-80	-61	-62	-9	-6	12	-9	-9	8	-16	-26	5	-32	-22	-5	-30	-26	-11	-17	-27	-14
BE	0	-14	-13	-54	-62	-50	-8	0	48	61	68	76	44	7	-30	-62	-71	-59	-84	-65	-64	-68	-37	-14	-29
CY	-2	-49	-86	-74	-44	-8	15	40	42	-25	-1	-46	-50	-50	-40	-41	-55	-25	-18	-22	-7	-8	-10	-20	-10
DE	1	0	-41	-35	-52	-78	-42	-28	17	19	13	58	21	-13	6	-13	-31	-59	-42	-25	-43	-32	-27	-21	-23
EE	-1	-53	-17	19	54	40	-3	20	-39	-69	-59	-83	-58	-45	-17	-10	-28	-21	-14	-18	-11	-11	0	-20	-14
ES	0	-21	-77	-16	80	134	127	24	-105	-171	-127	-85	-71	-53	-36	-16	-9	-28	-11	-4	-9	-14	-3	-4	-1
FR	-4	-75	-76	-26	18	-2	25	-11	-17	-26	-53	-32	-32	-41	-28	-32	-28	-26	-22	-12	-13	-27	-17	-3	-6
GR	0	-29	-65	-64	-43	-38	-12	21	1	43	11	37	21	-9	-38	-38	-40	-32	-44	-62	-44	-40	-32	-23	-19
IE	0	-19	-57	-56	-42	-19	20	66	83	61	65	12	-17	-17	-69	-85	-78	-39	-77	-48	-40	-53	-42	-32	-24
IT	0	-41	-58	-27	-34	27	8	29	-1	-49	-48	-17	-8	-40	-50	-13	-29	-22	-28	-3	-18	-35	-10	-17	-20
LU	0	-28	-65	-61	-29	-8	-6	11	-3	-23	-9	-31	-19	-27	-47	-23	-13	-36	-32	-33	-18	-22	-26	-25	-11
NL	0	-6	-33	-39	-49	-35	18	56	74	208	230	179	57	12	-26	-61	-91	-117	-107	-86	-73	-59	-61	-32	-31
PL	-11	-75	-59	-58	-21	6	17	11	15	-8	-7	-3	-32	-48	-19	-12	-40	-27	-36	-16	-23	-6	-36	-36	-26
PT	0	-5	-28	-17	-34	-4	58	57	168	182	180	70	27	-196	-169	-136	-106	-37	-57	-29	-8	-16	-9	-16	-16
SK	0	-50	-74	-42	8	21	13	18	6	-2	-19	-15	-39	-77	-26	-38	-45	-33	-29	-19	-19	-21	-12	-14	-12

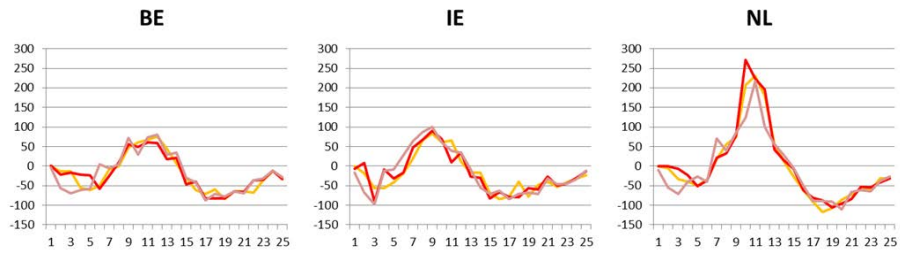
Proportional, linear on wealth	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
AT	-1	-21	-57	-17	-102	-57	-62	-9	10	0	-6	-7	31	-24	-12	-14	-19	-30	-6	-55	-12	-43	-22	-24	-12
BE	2	-22	-17	-22	-24	-58	-25	9	55	49	60	59	18	21	-47	-39	-82	-82	-83	-64	-66	-37	-35	-11	-33
CY	-15	-36	-124	-96	-28	17	38	68	10	-52	-18	-61	-25	-47	-42	-37	-42	-13	-24	-20	-11	-11	-3	-15	-19
DE	1	-12	-14	-36	15	-77	10	-49	18	14	4	24	15	-3	-15	-10	-43	-43	-15	-30	-53	-4	-30	-15	-14
EE	-59	-94	-42	25	113	105	12	-16	-16	-63	-61	-77	-52	-58	-5	-15	-24	-22	-28	-16	-13	-18	6	-8	-8
ES	-17	-12	-82	2	132	216	115	58	-54	-227	-130	-59	-73	-46	-15	-12	-20	-32	-5	-10	-8	-19	-9	3	-9
FR	-24	-15	-73	19	-17	31	39	-28	-28	-23	-55	-45	-60	-10	-43	-36	-28	-31	-14	-6	-20	-29	-3	-8	-4
GR	-11	-57	-66	-37	-63	-35	41	10	59	52	21	-5	13	-16	-36	-25	-36	-74	-33	-48	-49	-22	-37	-24	-3
IE	-7	8	-91	-9	-32	-16	48	67	90	69	10	35	-26	-30	-82	-66	-80	-79	-57	-59	-26	-49	-45	-30	-14
IT	-19	-79	-6	0	-53	-17	55	-6	-19	-16	-19	-35	-6	-30	-39	-17	-24	-22	-24	-19	-16	-34	-23	-29	-10
LU	3	-25	-28	-41	-28	-34	26	13	-40	-17	-18	-48	-56	-55	-27	-20	-10	-40	-14	-24	-23	-18	-43	-26	-24
NL	0	-1	-7	-24	-51	-36	21	32	77	271	223	195	42	15	-7	-61	-81	-87	-106	-96	-85	-53	-55	-41	-31
PL	-75	-158	-76	10	27	13	63	31	57	9	-21	8	-30	-34	-43	-17	-35	-17	-30	-30	-14	-19	-37	-33	-23
PT	-4	1	-6	-22	-45	-14	72	42	187	182	140	58	23	-164	-151	-104	-87	-33	-36	-24	-18	-16	0	-15	-15
SK	-65	-208	-95	54	43	42	84	67	14	21	-47	-5	-47	-26	-40	-27	-33	-44	-30	-16	-26	-4	-3	-17	-9

Proportional, median-centered	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
AT	-20	-146	-143	-45	-2	22	5	55	43	16	0	8	36	-23	-11	-14	-15	-18	0	-69	-3	-40	-25	-34	-14
BE	-3	-57	-69	-62	-58	4	-6	5	73	30	74	81	30	34	-30	-42	-87	-71	-78	-65	-70	-37	-32	-12	-29
CY	-59	-158	-83	-50	32	68	65	86	25	-46	-29	-48	-25	-42	-49	-19	-56	-11	-25	-26	-19	-14	2	-20	-18
DE	-4	-85	-68	-112	-28	-38	75	1	33	42	-22	23	26	3	-6	27	-42	-52	-27	-25	-44	1	-30	-12	-11
EE	-75	-94	61	41	121	126	3	-7	-35	-76	-40	-65	-69	-46	-14	-19	-28	-14	-17	-15	-12	-15	4	-7	-13
ES	-30	-78	-77	34	161	163	133	47	-55	-215	-133	-59	-72	-46	-16	-21	-26	-31	-6	-12	-7	-21	-8	1	-5
FR	-98	-123	3	109	34	38	73	-27	-29	-21	-40	-36	-49	-17	-40	-37	-23	-14	-24	-14	-12	-22	-3	-11	-2
GR	-20	-142	-94	-40	-27	-33	48	38	48	87	31	20	9	-26	-32	-30	-11	-62	-31	-40	-31	-22	-51	-18	-8
IE	-16	-68	-97	-11	-8	28	64	87	101	60	39	35	-9	-54	-71	-63	-85	-71	-68	-71	-31	-46	-44	-34	-11
IT	-50	-180	-40	89	-3	29	36	23	6	-42	-15	-33	1	-40	-57	-6	-25	-12	-26	-13	-12	-33	-21	-25	-12
LU	-29	-179	-58	12	54	30	48	21	-6	8	-3	-34	-58	-52	-26	-20	-13	-36	-19	-33	-18	-14	-34	-26	-17
NL	-10	-55	-71	-39	-27	-40	70	39	85	124	217	101	56	28	-6	-50	-89	-90	-91	-110	-66	-62	-64	-38	-27
PL	-95	-207	-33	70	49	33	53	18	47	0	-35	-6	-43	-27	-39	-21	-35	-21	-23	-34	-29	-18	-41	-31	-18
PT	-5	-33	-30	-43	-29	1	82	51	158	201	143	72	19	-168	-157	-112	-90	-30	-40	-27	-19	-10	-5	-15	-18
SK	-65	-157	-8	69	35	37	68	32	-8	1	-66	-5	-47	-49	-21	-31	-37	-39	-27	-20	-29	-4	-2	-12	-3

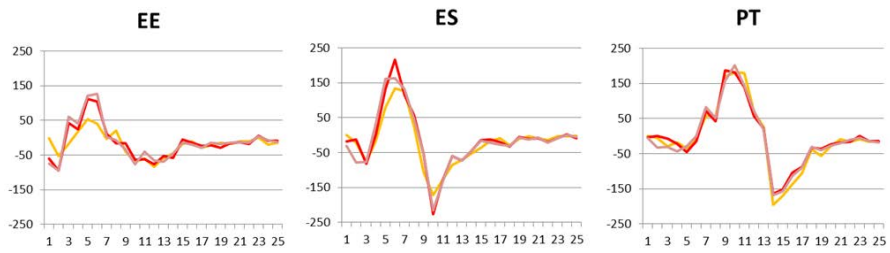
These tables show the absolute change in traded properties for all three caps cases with respect the no cap case.

We also present the information in the table above by country, grouped into the five categories described in Section 3.2.

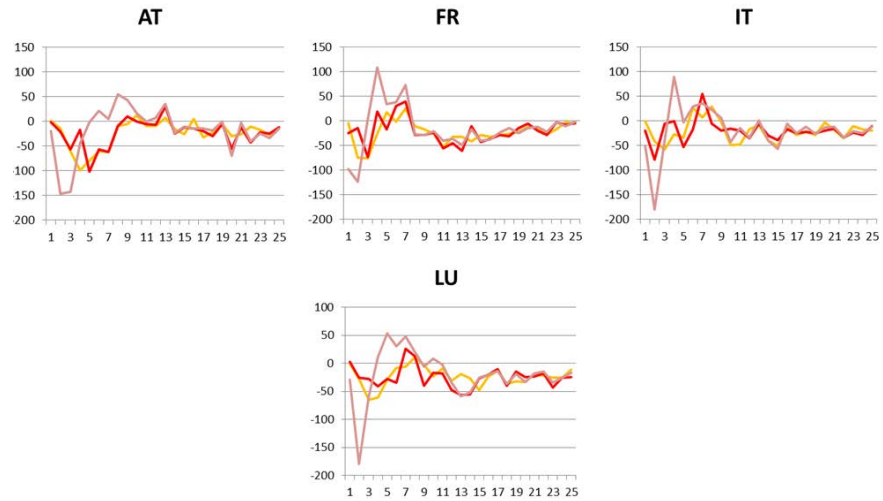
Group 1. "Price-symmetric" countries



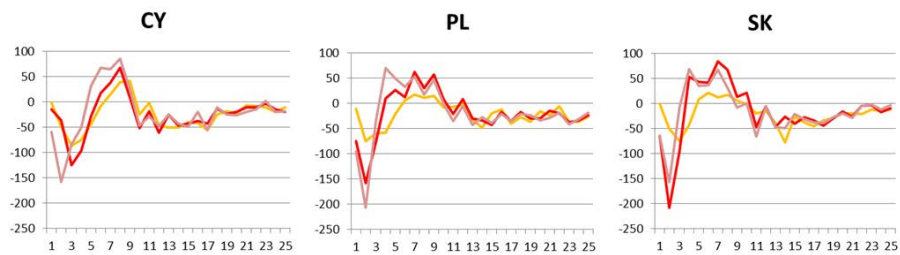
Group 2. "Wannabe-targeting" countries




Group 3. "Subprime-targeting" countries



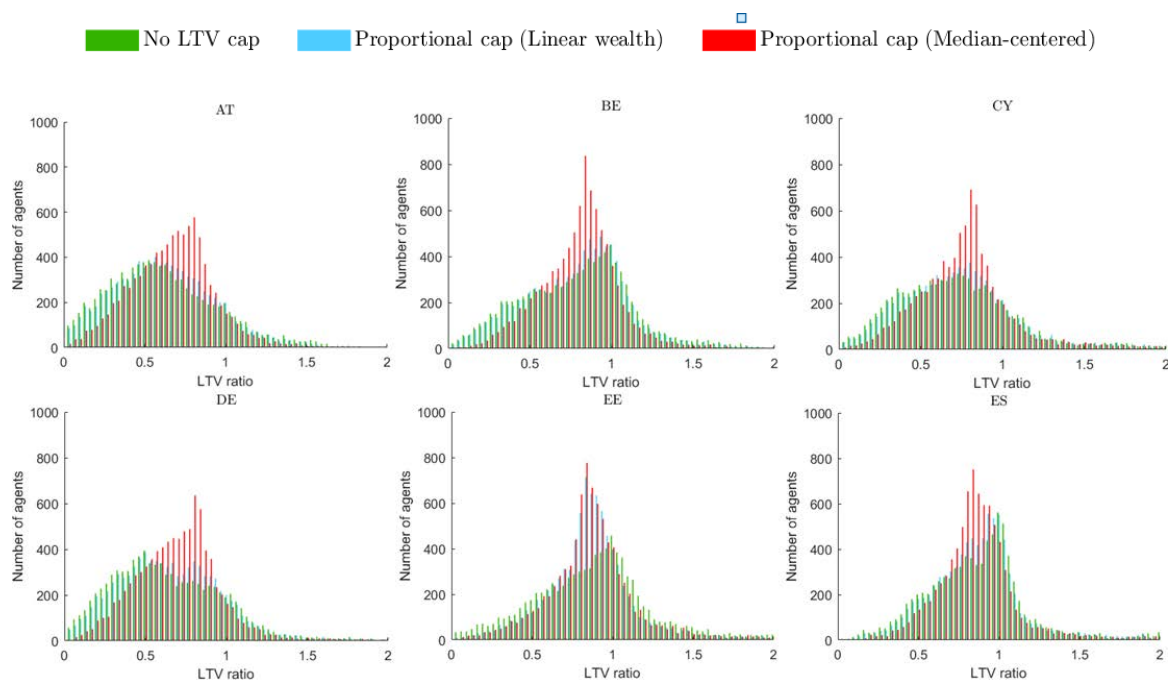
Group 4. "Flat" countries




Group 5. Other countries



Finally, the following matrix of graphs (Figure A3.2) allows to gauge the effect of the two proportionate LTV caps in the initial loan-to-value distributions. We do not show the histograms in the case of the absolute cap for the sake of clarity, as they usually entail a concentration of buyers around the cap value combined with an abrupt truncation of the distribution.



■ No LTV cap
 ■ Proportional cap (Linear wealth)
 ■ Proportional cap (Median-centered)



In this figure we show the effect of the two proportionate LTV caps in the initial loan-to-value distributions.

Acknowledgements

This paper should not be reported as representing the views of the European Central Bank, Banco de España or the International Monetary Fund. We thank Marco Gross, participants at 43rd Annual Conference of the Eastern Economic Association, participants at SCE 23rd International Conference on Computing in Economics and Finance and anonymous referees. We also acknowledge the financial support of Junta de Castilla-La Mancha grant PEI11- 0031-6939 and the Ministerio de Educación grant ECO2017-89715-P (Javier Población). Any errors are entirely the authors'.

Dimitrios Laliotis (corresponding author)

European Central Bank, Frankfurt am Main, Germany; International Monetary Fund, Washington, D.C., United States;
email: dlaliotis@imf.org

Alejandro Buesa

European Central Bank, Frankfurt am Main, Germany; Universidad Complutense de Madrid, Madrid, Spain; email: abuesa@ucm.es

Miha Leber

European Central Bank, Frankfurt am Main, Germany; email: miha.leber@ecb.europa.eu

Francisco Javier Población García

European Central Bank, Frankfurt am Main, Germany; email: javier.poblacion@bde.es

© European Central Bank, 2019

Postal address 60640 Frankfurt am Main, Germany

Telephone +49 69 1344 0

Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the [Social Science Research Network electronic library](#) or from [RePEc: Research Papers in Economics](#). Information on all of the papers published in the ECB Working Paper Series can be found on the [ECB's website](#).

PDF

ISBN 978-92-899-3556-2

ISSN 1725-2806

doi:10.2866/71742

QB-AR-19-075-EN-N