CAHIER D'ÉTUDES
WORKING PAPER

N° 161

USING HOUSEHOLD-LEVEL DATA TO GUIDE BORROWER-BASED MACRO-PRUDENTIAL POLICY

GASTON GIORDANA

MICHAEL ZIEGELMEYER

JULY 2022

BANQUE CENTRALE DU LUXEMBOURG

EUROSYSTÈME

Using household-level data to guide borrower-based macro-prudential policy[♣]

Gaston Giordana[♦] and Michael Ziegelmeyer[♠]

05 July 2022

Abstract

In 2019, Luxembourg introduced borrower-based instruments in the macro-prudential toolkit to constrain credit to households who exceed a certain limit on their loan-to-value ratio, on their (mortgage) debt-to-income ratio or on their debt service-to-income ratio. This paper analyses the impact of setting these limits at different levels, using household-level data from Luxembourg. We calculate these debt burden ratios for individual households who recently purchased their main residence using data from the Household Finance and Consumption Survey conducted in 2010, 2014 and 2018. On January 1, 2021 authorities imposed a legally binding limit on the loan-to-value (LTV) ratio for new mortgages. This may be 80%, 90% or 100% depending on the category of borrower. Had the least restrictive LTV limit envisaged by the law (100%) been applied in 2018, credit would have been rationed to 24% of households with recent mortgages on their main residence. This limit would have required a 7% reduction in the overall debt of this group of households. Had the most restrictive LTV limit envisaged by the law (75%) been applied in 2018, credit would have been rationed to 64% of households with recent mortgages on their main residence, requiring an 18% reduction in overall debt in this group. More generally, we evaluate how well borrower-based instruments can target those households that are financially vulnerable (according to conventional measures from the literature). By simulating an adverse scenario, we find that combining several ratios one could better target households that were not financially vulnerable in the benign conditions of 2018 but would become vulnerable after a shock to income. However, any borrower-based instrument inevitably generates some classification errors (either granting credit to households that are financially vulnerable or constraining credit to households that are not financially vulnerable). Using different assumptions on policymaker preferences, we apply the signals approach to derive limits that are "optimal" in the sense of minimising classification errors.

JEL-codes: D10, D14, G21, G28

Keywords: Household debt; Financial vulnerability; Macro-prudential policy

---

[♣] This paper should not be reported as representing the views of the BCL or the Eurosystem. The views expressed are those of the authors and may not be shared by other research staff or policymakers in the BCL, the Eurosystem or the Eurosystem Household Finance and Consumption Network. We give special thanks to Paolo Guarda. In addition, we would like to thank Giuseppe Pulina, Emmanuel Thibault, Martial Dupaigne, our colleagues from the Financial Stability Department and internal seminar participants for useful comments.

[♦] Economics and Research Department, Banque centrale du Luxembourg. E-mail: gaston_andres.giordana@bcl.lu, tel. +352 4774 4553, fax +352 4774 4920.

[♠] Economics and Research Department, Banque centrale du Luxembourg. E-mail: michael.ziegelmeyer@bcl.lu, tel. +352 4774 4566, fax +352 4774 4920.

Non-technical summary

This paper analyses borrower-based macro-prudential policy in Luxembourg using anonymous household-level data, providing a complementary approach to studies that use bank-level data or aggregate data. First, we calculate debt burden ratios and financial vulnerability measures for individual households. We compare results using data from the most recent wave of the Luxembourg Household Finance and Consumption Survey to those from previous survey waves in 2014 and 2010. Second, we simulate the effect of the borrower-based macro-prudential instruments passed into law in 2019, calculating how different policy settings would have affected mortgage credit to households with recent mortgages in 2018, as well as the volume of credit that would have exceeded regulatory limits. Third, we evaluate how well borrower-based measures can target households that are financially vulnerable (according to conventional measures from the literature). Lower settings of borrower-based measures classify more households as financially vulnerable, reducing the amount of credit available. Higher settings identify fewer households as financially vulnerable, allowing more credit but carrying potential long-term costs in terms of higher systemic risk. Policymakers cannot avoid this trade-off between different kinds of classification errors (either granting credit to households that are financially vulnerable or constraining credit to households that are not financially vulnerable). Using different assumptions on policymaker preferences, we apply a data-driven approach to derive the settings that are "optimal" in the sense of minimising classification errors.

The analysis leads to three main results. First, in Luxembourg the median loan-to-value ratio and debt-to-income ratio increased between 2014 and 2018. However, differences are not statistically significant, so we find no evidence that the household debt burden worsened dramatically. Using conventional definitions of financial vulnerability, the share of vulnerable households did not change significantly.

Second, for the loan-to-value ratio the **least restrictive** limit envisaged by the law (100%) would have affected 24% of households who borrowed to buy their main residence between 2015 and 2018. Affected households would have accounted for 28% of overall debt in this group, requiring a 7% reduction in their debt to comply with the limit. At the other extreme, the **most restrictive** limit on the loan-to-value ratio envisaged by the law (75%) would have affected 64% of households with recent mortgages, who represent 74% of overall debt in this group and would have had to reduce their debt by 18% to comply with the limit. Some of the limits the law envisages for other ratios are even more restrictive. A 400% limit on the ratio of mortgage debt to disposable income would have affected 76% of households with recent mortgages, who accounted for 87% of debt in this group and would have had to reduce their debt by 41.5% to comply with the limit.

Third, by simulating an adverse scenario, we show that simultaneously considering several ratios would be more effective at identifying households who were not financially vulnerable in the benign conditions of 2018, but would become vulnerable after a severe income shock. This suggests that a regulatory limit on the loan-to-value ratio complements rather than substitutes for banks' internal practices in screening potential borrowers.

Résumé non-technique

Cet article analyse les mesures macroprudentielles appliquées aux crédits immobiliers résidentiels au Luxembourg en utilisant des données anonymes au niveau des ménages individuels. Ainsi, il fournit une approche complémentaire aux études qui se basent sur les données au niveau des banques ou les données agrégées. Premièrement, nous calculons différents ratios d'endettement et mesurons la vulnérabilité financière des ménages individuels en utilisant les données de la plus récente vague de l'enquête luxembourgeoise sur le comportement financier et de consommation des ménages (collectée en 2018). Nous comparons ces résultats aux vagues précédentes, collectées en 2014 et 2010. Deuxièmement, nous simulons l'effet des différentes mesures macroprudentielles prévues par la loi du 4 décembre 2019, en calculant l'impact que différents paramétrages auraient pu avoir sur les ménages ayant un prêt hypothécaire récent en 2018, ainsi que le montant de crédit qui aurait dépassé les limites réglementaires. Troisièmement, nous évaluons la précision des instruments macroprudentiels pour cibler les ménages qui sont financièrement vulnérables (selon des critères typiques de la littérature). En fixant les limites macroprudentielles plus bas, davantage de ménages sont classés comme financièrement vulnérables, ce qui réduit le montant du crédit disponible. En fixant les limites plus haut, moins de ménages sont classés comme financièrement vulnérables, ce qui permet l'octroi de plus de crédit mais implique un risque systémique plus élevé avec des coûts potentiels à long terme. Les autorités ne peuvent pas éviter cet arbitrage entre différents types d'erreur (l'allocation de crédit à des ménages vulnérables ou la restriction de crédits à des ménages qui ne sont pas vulnérables). En utilisant différentes hypothèses quant aux préférences des autorités, nous appliquons une approche guidée par les données pour calculer le paramétrage « optimal », dans le sens qu'il minimise les erreurs de classification.

L'analyse met en exergue trois résultats principaux. Premièrement, entre 2014 et 2018 la valeur médiane du ratio prêt/valeur a augmenté au Luxembourg, comme aussi celle du ratio dette-sur-revenu disponible. Cependant, les différences ne sont pas statistiquement significatives et donc on ne peut pas affirmer que la charge de la dette des ménages s'est détériorée considérablement. En utilisant des définitions conventionnelles de la vulnérabilité financière, la part des ménages vulnérables n'a pas changé de manière significative.

Deuxièmement, pour le ratio prêt/valeur (*loan-to-value ratio*, LTV), la limite la **moins restrictive** qui est envisagée par la loi (100%) aurait affecté 24% des ménages qui ont emprunté entre 2015 et 2018 pour acheter leur résidence principale. Les ménages concernés auraient représenté 28% de la dette totale de ce groupe et auraient dû réduire leur dette de 7% pour se conformer à la limite. Par contre, la limite la **plus restrictive** qui est envisagée par la loi (75%) aurait affecté 64% des ménages avec des hypothèques récentes et 74% de la dette totale de ce groupe, exigeant une réduction de 18% de leur dette pour se conformer à la limite. La loi envisage des limites pour d'autres ratios qui peuvent être encore plus restrictives. Une limite de 400% pour le ratio de la dette hypothécaire au revenu disponible aurait affecté 76% des ménages avec des prêts hypothécaires récents et 87% de la dette de ce groupe, exigeant une réduction de 41,5% de la dette des ménages concernés.

Troisièmement, en simulant un scénario défavorable, nous montrons que la combinaison de plusieurs ratios est plus efficace pour identifier les ménages qui n'étaient pas financièrement vulnérables dans les conditions bénignes de 2018, mais qui deviendraient vulnérables après un choc sévère sur leur

revenu. Cela suggère qu'une limite réglementaire sur le ratio prêt/valeur est plus un complément qu'un substitut des pratiques internes appliquées par les banques lors de la sélection des emprunteurs.

1 Introduction

In 2019, Luxembourg extended its macro-prudential policy toolkit to introduce borrower-based instruments. These allow authorities to constrain the supply of credit for house purchase by setting an upper limit on various measures of a borrower's debt burden, including the loan-to-value ratio, the debt-to-disposable income ratio, or the debt service-to-disposable income ratio. In November 2020, Luxembourg's Systemic Risk Committee issued a recommendation[1] to impose legally binding limits on the loan-to-value ratio in the residential property market. A month later, the regulatory authority issued a regulation[2] requiring credit institutions to keep the loan-to-value ratio below 80% for most residential mortgage loans (including the "buy-to-let" segment), below 90% for households buying their main residence[3] and below 100% for first-time buyers acquiring their main residence.

This paper uses data from the most recent wave of the Luxembourg Household Finance and Consumption Survey (LU-HFCS) to update simple household debt burden ratios and more complex measures of financial vulnerability last published in Giordana and Ziegelmeyer (2017, 2020). Our measures of financial vulnerability draw on work by Ampudia et al. (2016) and by Meriküll and Rõõm (2020).[4] We then perform an ex-ante evaluation of the effect of various borrower-based macroprudential instruments on household financial vulnerability. Following Albacete et al. (2018) and Bańbula et al. (2016), we adapt the signals approach in Detken et al. (2014) to identify the settings of borrower-based instruments that most effectively target financially vulnerable households. Low settings will classify more households as financially vulnerable, carrying short-term costs in terms of lower economic activity and welfare. However, high settings will identify fewer households as financially vulnerable, generating long-term costs in terms of higher systemic risk. Our study is innovative in considering combinations of several debt burden ratios with households being only credit constrained if they exceed several limits simultaneously. Moreover, our study illustrates how household-level data can help to set borrower-based instruments despite potential caveats discussed below. Thus, it complements more aggregate approaches using macro-economic models, such as Sangaré (2019), or approaches based on bank data, which may contain less complete and less harmonized information about individual borrowers.

The 2019 law[5] allows authorities to impose limits on various ratios measuring a borrower's debt burden. We calculate how different values on these limits would have affected households with recent mortgages on their main residence, as well as the amount of credit that would not have been granted. Luxembourg authorities imposed a limit on the loan-to-value (LTV) ratio, where the least restrictive limit envisaged by the law (100%) would have affected 24% of households with recent mortgages on their main residence. This would have required a 7% reduction in overall debt of households with recent mortgages on their main residence to comply with the limit. The most restrictive LTV limit envisaged by the law (75%) would have affected 64% of households with recent mortgages. This would

1 See http://cdrs.lu/9-novembre-2020/

2 See http://data.legilux.public.lu/file/eli-etat-leg-rcsf-2020-12-03-a969-jo-fr-pdf.pdf

3 In the category of loans for principal residences granted to borrowers who are not first-time buyers, lending institutions may grant individual loans with an LTV ratio of up to 100%, provided that the aggregate amount of the loans benefiting from this derogation (90% < LTV ratio ≤ 100%) represents no more than 15% of the annual aggregate amount of this category of loans granted by this institution.

4 Leika and Marchettini (2017) propose a generalised framework to analyse financial vulnerability.

5 See http://data.legilux.public.lu/file/eli-etat-leg-loi-2019-12-04-a811-jo-fr-pdf.pdf

have required an 18% reduction in overall debt of households with recent mortgages on their main residence.

Borrower-based instruments aim to increase resilience and prevent financial crises. Therefore, they should not just focus on households that are financially vulnerable today, but also on those that may become financially vulnerable under adverse economic conditions. We therefore implement the stress test framework in Giordana and Ziegelmeyer (2020) to simulate household balance sheets under adverse conditions. We compare the benign economic environment in 2018 to an adverse economic scenario featuring a household income shock consistent with a substantial increase in unemployment, comparable to those the Covid-19 pandemic triggered in several European countries. Using a data-driven approach to evaluate alternative policies, we find that considering several ratios simultaneously is more effective to identify households who are not financially vulnerable in the baseline but become vulnerable following the income shock.[6] Among the policies we consider, the best performance requires restricting credit to households that breach limits on at least three out of four ratios. This suggests regulatory limits on the loan-to-value ratio should be considered complementary to banks' internal practices in screening potential borrowers and that the effectiveness of regulatory LTV limits could be increased if complemented by additional regulatory limits on other ratios.

Although ex-ante evaluations of borrower-based instruments are common (CGFS, 2016), their results inevitably depend on the underlying assumptions. In particular, our results are subject to several caveats. First, we do not observe household defaults, since there is no functioning credit register in Luxembourg. Therefore, we use survey data to identify financially vulnerable households by comparing their financial margin and liquid assets. While this is a limitation when monitoring developments in household credit risk, it is an advantage when simulating counterfactual scenarios for policy evaluation. Second, our financial vulnerability measure requires an estimate of disposable income and basic living costs, which introduces some uncertainty in the results. Third, our sample is limited to households who were actually granted mortgages, ignoring those who were refused credit and remained renters. Since households who rent have lower income[7], this may introduce a selection bias, possibly lowering our estimates of "optimal" policy settings. However, this bias should be negligible if the number of rejections is limited or if rejections are correlated with our measure of financial vulnerability, which should anyway resemble the internal indicators lenders use to screen potential borrowers. Moreover, our limited sample size increases the uncertainty around our point estimates, in particular for the optimal policy settings, which should be interpreted under this caveat. Fourth, our approach does not consider possible non-linear effects on financial stability if vulnerable households are concentrated in systemically important banks, rather than distributed evenly across lenders. Given the high concentration of the residential mortgage market among a few domestic banks in Luxembourg, this effect could be important. Fifth, we do not consider how private agents may respond to the introduction of a borrower-based measure. Single-income families and those in part time work may increase labour supply and therefore household income. Households may also reduce consumption, affecting their financial margin and therefore their probability of default. In addition, households may reduce their demand for mortgages and house prices may rise more slowly, affecting

6 Other studies reached similar conclusion using different methodologies, data and approaches. Grodecka's (2017) theoretical exercise provides the arguments underlying our results. In a DSGE framework calibrated for Sweden, Chen and Columba (2016) find that a mix of macroprudential measures is needed to deliver the maximum welfare benefit.

7 Around one-third of households in Luxembourg rent their household main residence (Chen et al., 2020). Compared to homeowners, renters have lower income and wealth (Chen et al., 2021).

the value of collateral and therefore the distribution of loss given default across the population of households. Finally, the signals approach does not account for household heterogeneity, as it performs a binary classification of households as either vulnerable or non-vulnerable, ignoring the extent of their financial vulnerability or the size of their mortgage. However, household heterogeneity in these dimensions is likely to affect short-run costs and long-term benefits, as we explore in a separate section.

Although we consider all potential borrower-based measures in the Luxembourg law, our analysis necessarily ignores many practical considerations that may have influenced policymakers. In particular, the loan-to-value (LTV) ratio may have several advantages compared to other borrower-based instruments. First, the LTV ratio may be more transparent, since it focuses on collateral in a financial transaction and is therefore directly observable from administrative data collected at the time of purchase. Instead, alternative measures such as the debt-to-income ratio or the debt service-to-income ratio may be more challenging to use because they require an estimate of household disposable income. While banks can observe movements on their clients' bank accounts over several months, borrowers may choose not to report their debt service obligations at other financial institutions (hence the need for a credit register). Second, the LTV ratio may be more comparable across lenders and households, who may be applying different definitions of household disposable income. Third, a regulatory limit on the LTV ratio may complement lenders' credit-scoring and screening practices, which rely on internal estimates of disposable income, debt service and outstanding debt. Fourth, the LTV ratio may have been politically more acceptable than ratios that refer to household disposable income, because the latter necessarily separate low-income households from high-income households. This may have been an important advantage since Luxembourg introduced LTV limits in the context of the pandemic, when there may have been concerns regarding the potential procyclical effects of income-based measures in a downturn. Therefore, given the underlying risks to the real estate sector in Luxembourg, the LTV ratio may have been preferred because it could better target these vulnerabilities, with fewer side effects for the rest of the economy. Finally, the LTV ratio has a direct effect on bank resilience, clearly falling in the scope of macro-prudential policy.

The paper is organized as follows. Section 2 introduces the data, presents the debt burden ratios as defined in the law and reports statistics on the debt burden constructed from the data. Section 3 introduces the measure of financial vulnerability that we draw from the literature. Sub-section 3.2 describes our adverse scenario and reports its impact on financial vulnerability. Sub-section 4.1 describes the signals approach that we use to evaluate alternative policies and derive "optimal" limits that minimise classification errors. Sub-section 4.2 reports the results, including the "optimal" limits in the baseline and in the adverse scenario. Section 5 analyses the impact of household heterogeneity on the costs and benefits of borrower-based measures. Section 6 discusses how our results can inform policy when setting limits on the loan-to-value ratio. Section 7 concludes.

2 Data and Methodology

2.1 Data

This paper uses household-level data from the first three waves of the LU-HFCS survey of households resident in Luxembourg. The first wave was conducted mostly in 2010 and included 950 households, the second wave was conducted in 2014 and included 1601 households, and the third wave was conducted in 2018 with 1616 households. Samples were designed to be representative of the entire population of households resident in Luxembourg and are weighted accordingly.

We focus on households that only recently took out their mortgage(s), to better estimate the impact from activating borrower-based instruments. Within this group of recent borrowers, we focus on mortgages on the household main residence (HMR), excluding recent mortgages on other real estate property (OREP).[8] Each survey wave requires a different definition of what qualifies as a "recent" mortgage. For the third wave, "recent mortgages" date between 2015 and 2018, for the second wave between 2011 and 2014, and for the first wave between 2007 and 2011. Table 1 provides the number of observations in each subgroup, as well as the corresponding share of the population and of overall outstanding debt (including non-mortgage debt).

The analysis below will focus on households with recent HMR mortgages, since this group is more homogeneous and would be directly affected by the introduction of borrower-based measures. The bottom row in Table 1 reports that households with recent HMR mortgages account for about a third of all household debt. In a robustness exercise reported in Table 13, we extend this group to include households with recent mortgages on OREP. Conclusions are unchanged, although this extended sample covers nearly half of all household debt (row before last in Table 1).

Table 1: Number of observations and share of the population represented by subgroup and survey wave

	unweighted number of obs			% of population (weighted)			% of outstanding debt (weighted)		
Indebted households	**2010**	**2014**	**2018**	**2010**	**2014**	**2018**	**2010**	**2014**	**2018**
All households with outstanding debt	580	952	960	58.3%	54.6%	53.2%	100%	100%	100%
Households with mortgage debt	405	664	640	38.8%	35.2%	31.2%	95%	94%	94%
Households with HMR mortgage	328	547	539	32.8%	29.1%	26.8%	79%	80%	78%
Households with recent mortgage debt	152	215	210	14.8%	11.5%	10.8%	49%	45%	46%
Households with recent HMR mortgage debt	111	149	144	11.4%	8.5%	7.8%	37%	33%	31%

Source: Own calculations based on wave 1, 2 and 3 of the LU-HFCS; data are multiply imputed.
Note: The survey refers to a population of 186,440 households in 2010, 210,965 households in 2014 and 226,378 households in 2018.
Overall debt includes all loans (mortgage and non-mortgage) from financial institutions or from relatives, friends, employers, etc.. Mortgage debt may be for the household main residence (HMR) or other real estate property (OREP). Household overall debt was €15,200 million in 2010, €20,500 million in 2014 and €23,700 million in 2018.

In general, survey data suffer from a bias due to underreporting and missing responses, especially among the wealthiest households. To limit this bias, HFCS data is multiply imputed and statistics

8 The HFCS questionnaire contains the following question on other real estate property: *"Apart from your house/apartment [for owners], do you/does your household own any (other) properties, such as houses, apartments, garages, offices, hotels, other commercial buildings, farms, land, etc.? PROBE: Please include properties both here in Luxembourg and elsewhere."*

reported below account for uncertainty due to sampling and imputation methods. Results are based on 1000 replicate weights and 5 multiply imputed implicates of the dataset.

2.2 Definition of household debt burden ratios

Since Luxembourg does not yet have a functioning credit register, survey data is an invaluable source of household-level data to study financial vulnerability. For individual households in the survey, we calculate the ratios defined in the law on borrower-based instruments[9]. However, in some cases we can only approximate the ratios as defined in the law. In particular, survey data usually refers to the date of the survey, while borrower-based instruments rely on ratios calculated when the household applies for a loan. Therefore, in some cases we need to estimate a household's debt burden ratio in the year when it took out its mortgage (as opposed to the year of the survey). This involves adjustments to both debt and income. For total debt, we add repayments on other debts[10] between the date of the most recent mortgage and the date of the survey. For disposable income, we adjust for aggregate wage growth over the same period. These adjustments are smaller for households who had "recently" taken out a mortgage when the survey took place, which is another reason why we limit the analysis to this group. For the household main residence (HMR), the survey explicitly asks for the loan-to-value ratio at mortgage origination, so this ratio does not need adjustment. However, for other real estate properties (OREP) the survey only collects the amount of mortgage outstanding at the time of the survey. This requires an adjustment, possibly injecting some error, so the sample is only extended to include households with recent OREP mortgages in the robustness exercise (Table 13). Since we focus on ratios as defined in the Luxembourg law, results reported below may differ from those in previous ECB and BCL publications (HFCN, 2013, 2016, 2020; Giordana and Ziegelmeyer, 2017).

Table 2 describes the debt burden ratios as defined in the law (second column) and as calculated from survey data (third column). The initial loan-to-value (LTV) ratio indicates how far the property value covers potential loan losses and therefore measures the lender's exposure to counterparty credit risk. The two debt-to-income (DI) ratios provide an indication of households' ability to service their debt from their disposable income. The first ratio (MDI) focuses on mortgage debt, while the second ratio (DI) considers overall debt. The debt service-to-income (DSI) ratio compares the flow of regular debt payments to the flow of disposable income at the time of mortgage origination. Finally, mortgage maturity (MM) is not strictly speaking a debt burden ratio, but puts a limit on the mortgage maturity at loan origination[11]. Combining mortgage maturity limits with limits on the other ratios could help restrain potential feedback from household debt to property prices[12].

The LTV ratio and MM are observed at loan origination and do not evolve as a household gradually pays back the loan. Unlike DSI or DI ratios, they do not require an estimate of the household situation at the time the loan was granted (in terms of disposable income or total outstanding debt). This makes

9 http://legilux.public.lu/eli/etat/leg/loi/2019/12/04/a811/jo

10 The initial amount of the most recent HMR mortgage is reported directly in the survey. If the household has more than one mortgage, we adjust older mortgages using a linear approximation based on the difference between the initial value of the older loan and its outstanding amount at the survey date.

11 To facilitate reading, "debt burden ratios" in the text also includes the limit on mortgage maturity.

12 Evidence in Ferreira (2018) suggests that such a feedback loop may have been active in Luxembourg.

LTV and MM the easiest ratios to measure, even using the HFCS survey data, giving them a clear advantage from a practical point of view.

Table 2: Household debt burden ratios

Debt burden ratio	Legal definition*	Calculation with HFCS data	Limit range[(a)]
Loan-to-value ratio of HMR (LTV)	Art. I, 2°(2)a) *"the ratio of the amount of all mortgage loans or tranches of mortgage loans guaranteed by the borrower at the time of loan origination and the value of the property at the same time."*	Initial amount of HMR mortgages divided by the self-reported acquisition value of the HMR.	75% - 100%
Mortgage debt-to-disposable income ratio (MDI)	Art. I, 2°(2)b) *"the ratio of the amount of all mortgage loans or tranches of mortgage loans guaranteed by the borrower at the time of loan origination and the borrower's total annual disposable income at the same time."*	Total outstanding mortgage debt[(b)] divided by annual household disposable income[(b)].	400% - 1200%
Debt-to-disposable income ratio (DI)	Art. I, 2°(2)c) *"the ratio of the borrower's total debt at the time of loan origination and the borrower's total annual disposable income at the same time."*	Total outstanding debt[(b)] divided by annual household disposable income[(b)].	400% - 1200%
Debt service-to-disposable income ratio (DSI)	Art. I, 2°(2)d) *"the ratio of the borrower's annual total debt service and the borrower's total annual disposable income at the time of loan origination."*	Monthly debt payments[(c)] divided by monthly disposable income[(b)].	35% - 75%
Mortgage maturity (MM)	Art. I, 2°(2)e) *"Initial maturity of the loan."*	Maximum initial maturity of the two most important mortgages	20 - 35 years

** Own translation of the legal text in French.*
MDI, DI and DSI ratios are only available for wave 2 and 3 because wave 1 did not collect disposable income. Disposable income is adjusted using growth in aggregate disposable income (B7.b in household sector accounts) between the year of take out and the year of the survey, correcting for growth in the number of private households in Luxembourg (Eurostat lfst_hhnhtych).
(a) From legal text. (b) At the time of the most recent mortgage take out.
(c) HFCS includes no data on debt service of credit lines/overdraft liabilities (set to zero). Debt service includes interest and principal repayment but excludes taxes, insurance and any other related fees. Payments for leasing contracts are also excluded. As Albacete and Lindner (2017), we assume constant debt service (no change between the time of loan origination and the time of the survey).

2.3 Debt burden ratios

For the different debt burden ratios, Table 3 presents the median value across households with recent HMR mortgages in each of the three survey waves. Results for 2018 currently represent the most up-to-date assessment using household level data. In the following, we discuss the results for each debt burden ratio in turn.

For the loan-to-value (LTV) ratio, the median across households with a recent HMR mortgage was 88.6% in 2018, 9 percentage points higher than in 2014 (although the difference is not statistically significant). The result is similar when extending the sample to include households with recent mortgages on other real estate property (Table 13 in Appendix A).

Table 12 in the appendix reports the median LTV ratio for various subgroups in 2018. Not surprisingly, the median LTV is higher in households with a younger reference person or in the lower net wealth quintiles, but these differences are not statistically significant.[13]

Table 3: Debt burden ratios at loan origination: households with recent HMR mortgages

	Wave 2010		Wave 2014		Wave 2018	
Debt burden ratios	Median	Std. Err.	Median	Std. Err.	Median	Std. Err.
Initial loan-to-value (LTV) ratio of HMR	92.8%	7.5%	80.0%	4.5%	88.6%	3.7%
Mortgage debt-to-disposable income (MDI) ratio	n.a.		580.1%	42.9%	643.6%	36.7%
Debt-to-disposable income (DI) ratio	n.a.		600.7%	43.1%	647.0%	38.1%
Debt service-to-disposable income (DSI) ratio	n.a.		35.2%	1.1%	34.5%	2.6%
Mortgage maturity (MM) in years	25.00	1.35	25.00	1.38	25.00	0.49

Source: Own calculations based on waves 1, 2 and 3 of the LU-HFCS; data are multiply imputed and weighted; variance estimation based on 1000 replicate weights. Results for the debt service-to-income ratio differ from those in Girshina, Mathä and Ziegelmeyer (2017) because leasing payments are now included. Income and debt are measured at date of latest HMR mortgage origination. n.a. = disposable income not available in the 2010 wave. Differences in median debt burden ratios across years are not statistically significant.

Panel (a) in Figure 1 plots the cumulative distribution of the LTV ratio across households with recent HMR mortgages in different survey years. The vertical lines at values 0.75 and 1 represent the lowest and highest LTV limits envisaged by the law on borrower-based instruments (see Table 2). What appears to be a discontinuity at the value 1.00 is actually a step in the cumulative distribution, because many observations are clustered at this value, although when these HMR mortgages were granted the law had not yet passed. In 2018, 76% of households with recent HMR mortgages had an LTV ratio of 1.00 or below. Despite some overlapping at low percentiles, the distribution in 2014 is clearly to the left of the distribution in 2018, suggesting there was an increase in the share of households with higher LTV ratios. However, the Kolmogorov–Smirnov test cannot reject the null hypothesis that the distributions are equal.

In panel (b) of Figure 1, the most restrictive LTV limit envisaged by the law (75%) would affect 64% of households with recent HMR mortgages. The debt of households affected represents 74% of the overall debt of households with recent HMR mortgages, with 18% of overall debt in this group exceeding this particular LTV limit at the household level. On the right side of panel (b), the least restrictive LTV limit envisaged by the law (100%) would affect 24% of households with recent HMR mortgages, who represent 28% of overall debt in this group, with 7% of overall debt in this group exceeding this LTV limit at the household level.

---

[13] For each debt burden indicator, we run a quantile regression (results not reported) to confirm that age and wealth are not statistically significant when accounting for other household characteristics (listed in Table 12).

Figure 1: Initial loan-to-value (LTV) ratio

(a) cumulative distribution by survey year

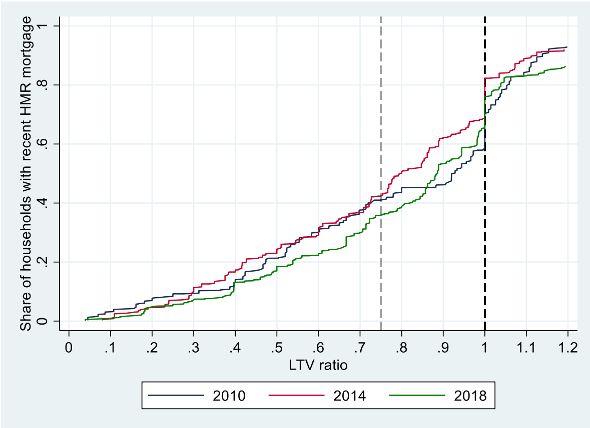

(b) borrowers affected in 2018, their share of overall debt and the share of overall debt exceeding limit

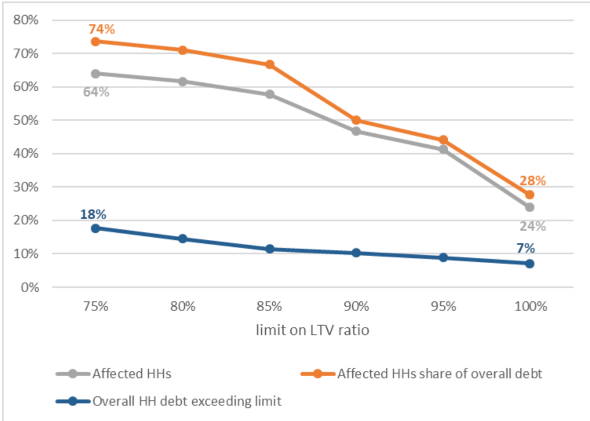

Source: Own calculations based on waves 1, 2 and 3 of the LU-HFCS; data are multiply imputed and weighted; only households with recent HMR mortgages.
Panel (a): cumulative distribution functions are calculated across all 5 implicates each year. Vertical lines indicate lowest and highest limits envisaged by the law. We omit the upper tail of the cumulative distribution functions.

For the mortgage debt-to-income (MDI) ratio, Table 3 above reported that the median in 2018 was 643.6%, 60 percentage points higher than in 2014, although the difference is not statistically significant. Extending the sample to include households with recent mortgages on other real estate property (Table 13 in Appendix A) also suggests that the median MDI ratio was higher in 2018, but the difference is still not statistically significant. According to Table 12 in the appendix, the median among households with a reference person aged between 45-54 years was 338.9%, significantly below the median for younger households. The median MDI ratio in 2018 was not statistically different across other household characteristics.

Panel (a) in Figure 2 seems to indicate an improvement at the top of the distribution. For MDI ratios above 9, the 2018 cumulative distribution is systematically above the 2014 distribution, suggesting that fewer households in 2018 had MDI ratios above these values. However, the Kolmogorov–Smirnov test cannot reject the null hypothesis of equal distributions.

Panel (a) of Figure 2 suggests that the range of possible limits on the MDI ratio is much broader than the range of possible limits on the LTV ratio. On the one hand, Panel (b) of Figure 2 suggests that more restrictive MDI limits are even more restrictive than more restrictive limits on the LTV ratio. The most restrictive MDI limit envisaged by the law (400%) would affect 76% of households with recent HMR mortgages, who represent 87% of total debt in this group, which would have to decline by 42% to comply with this specific MDI limit. On the other hand, the less restrictive MDI limits have substantially lower impacts than the least restrictive LTV limit. Thus, the least restrictive MDI limit envisaged by the law (1200%) would affect only 9% of households with recent HMR mortgages, who represent 14% of total debt in this group, with only 5% of overall debt in this group exceeding this particular MDI limit.

Figure 2: Mortgage debt-to-income (MDI) ratio

(a) cumulative distribution by survey year

(b) borrowers affected in 2018, their share of overall debt and the share of overall debt exceeding limit

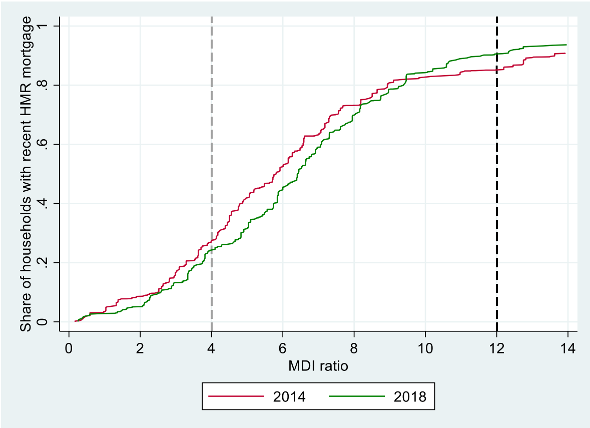

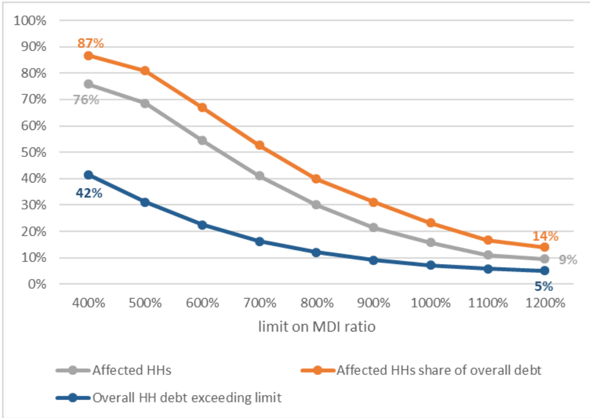

Source: Own calculations based on waves 2 and 3 of the LU-HFCS; data are multiply imputed and weighted; only households with recent HMR mortgages.
Panel (a): cumulative distribution functions are calculated across all 5 implicates for each year. Vertical lines indicate lowest and highest limits envisaged by the law. We omit the upper tail of the cumulative distribution functions.

Turning to the debt-to-income (DI) ratio, Table 3 above reported that the median DI ratio increased from 600.7% in 2014 to 647% in 2018 (the difference is not statistically significant). Comparing median DI ratios across groups with different household characteristics, results are similar to those for the MDI ratio (see Table 12 in the appendix).

Figure 3: Debt-to-income (DI) ratio

(a) cumulative distribution by survey year

(b) borrowers affected in 2018, their share of overall debt and the share of overall debt exceeding limit

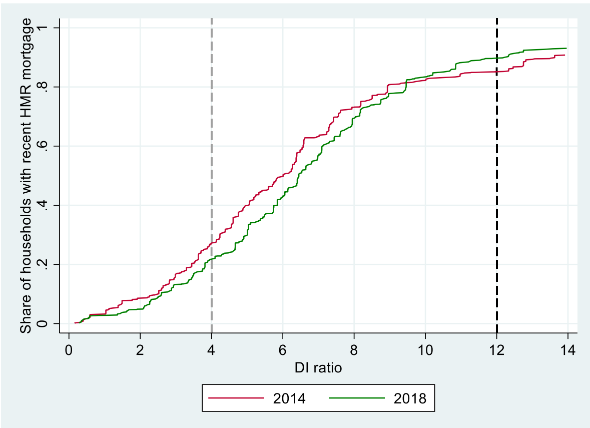

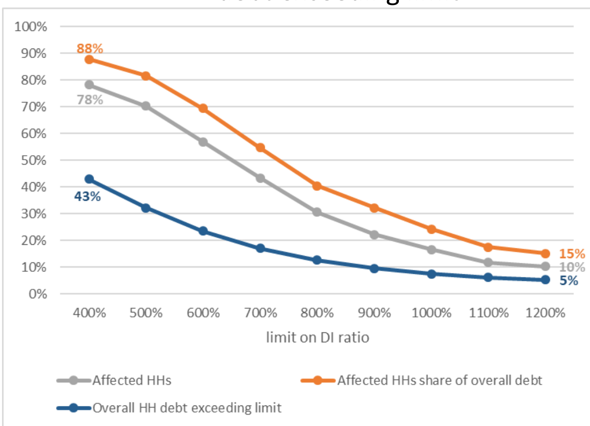

Source: Own calculations based on waves 2 and 3 of the LU-HFCS; data are multiply imputed and weighted; only households with recent HMR mortgages.
Panel (a): cumulative distribution functions are calculated across all 5 implicates for each year. Vertical lines indicate lowest and highest limits envisaged by the law. We omit the upper tail of the cumulative distribution functions.

Panel (a) in Figure 3 also suggests a possible improvement at the top of the distribution, as the cumulative distribution in 2018 rises above the 2014 distribution around a DI ratio of 9. This suggests that fewer households in 2018 had a DI ratio above this level. However, the Kolmogorov–Smirnov test again fails to reject the null hypothesis of equal distributions.

Panel (b) in Figure 3 indicates that the effects of limits on the DI ratio are very similar to those of limits on the MDI ratio. For example, the most restrictive DI limit envisaged by the law (400%) would affect 78% of households with recent HMR mortgages, who represent 88% of all overall debt in this group, which would have to decline by 43% to comply with this DI limit. The least restrictive DI limit envisaged by the law (1200%) would affect 10% of households with recent HMR mortgages, who represent 15% of overall debt in this group, with 5% of this overall debt exceeding this particular DI limit at the household level. The similarity with the effects of the MDI ratio are not surprising, given that for homeowners mortgage debt is the dominant component of overall household debt.

The debt service-to-income (DSI) ratio compares the flow of monthly debt payments to the flow of monthly gross income. Using the weighted data from the survey, the median interest rate on HMR mortgages declined from 2.20% in 2014 to 1.86% in 2018[14]. In Table 3, the median DSI ratio declined slightly to 34.5%in 2018 from 35.2% in 2014, but the difference is not statistically significant. Table 12 also finds no statistically significant differences in median DSI across groups by household characteristics. In Panel (a) of Figure 4, the Kolmogorov–Smirnov test again failed to reject the null hypothesis of equal distributions.

Figure 4: Debt service-to-income (DSI) ratio

(a) cumulative distribution by survey year

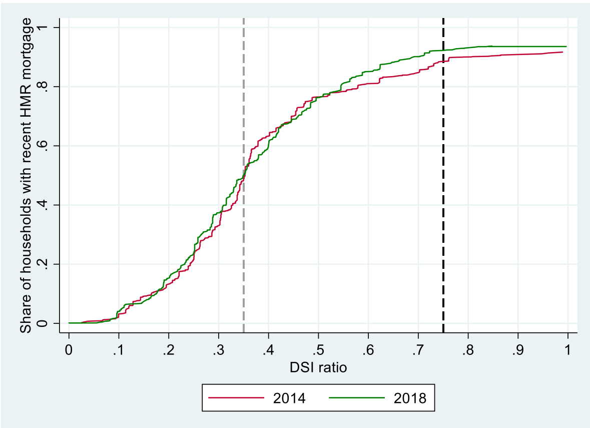

(b) borrowers affected in 2018, their share of overall debt and the share of overall debt exceeding limit

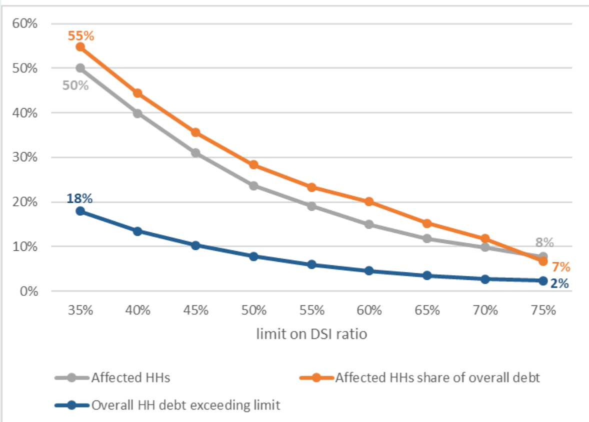

Source: Own calculations based on waves 2 and 3 of the LU-HFCS; data are multiply imputed and weighted; only households with recent HMR mortgages.
Panel (a): cumulative distribution functions are calculated across all 5 implicates for each year. Vertical lines indicate lowest and highest limits envisaged by the law. We omit the upper tail of the cumulative distribution functions.

Panel (b) of Figure 4 reveals that the most restrictive DSI limit envisaged by the law (35%) is very close to the median DSI of households with recent mortgages. Such a limit would affect 50% of households with recent HMR mortgages, who account for 55% of all total debt in this group, and would require an 18% reduction of overall debt in this group. Thus, restrictive DSI limits appear to be less restrictive in

[14] Mortgage rates reported by banks in Table 03.02 on www.bcl.lu declined from 2.15% in 2014 to 1.85% in 2018.

terms of the number of households affected, although they are as restrictive as LTV limits in terms of the volume of debt that would be affected.

Finally, the bottom row of Table 3 reports that the median maturity at mortgage origination was 25 years across the different waves of the survey. The estimate of median maturity was robust to extending the sample to include households with recent OREP mortgages (Table 13 in Appendix A).

For mortgage maturity at loan origination, Panel (a) in Figure 5 shows that the cumulative distribution follows a step function, as banks and borrowers usually agree on an initial mortgage maturity that is a multiple of five years.[15] The Kolmogorov-Smirnov test cannot reject the null hypothesis of equal distributions.

Panel (b) in Figure 5 indicates that the most restrictive maturity limit envisaged by the law (20 years) would affect 62% of households with recent HMR mortgages, representing 72% of all total debt in this group. The least restrictive maturity limit (35 years) would affect no households with recent HMR mortgages, as no households in the sample reported HMR mortgage maturities longer than 30 years.

Figure 5: Mortgage maturity at origination

(a) cumulative distribution by survey year

(b) borrowers affected in 2018 and their share of overall debt

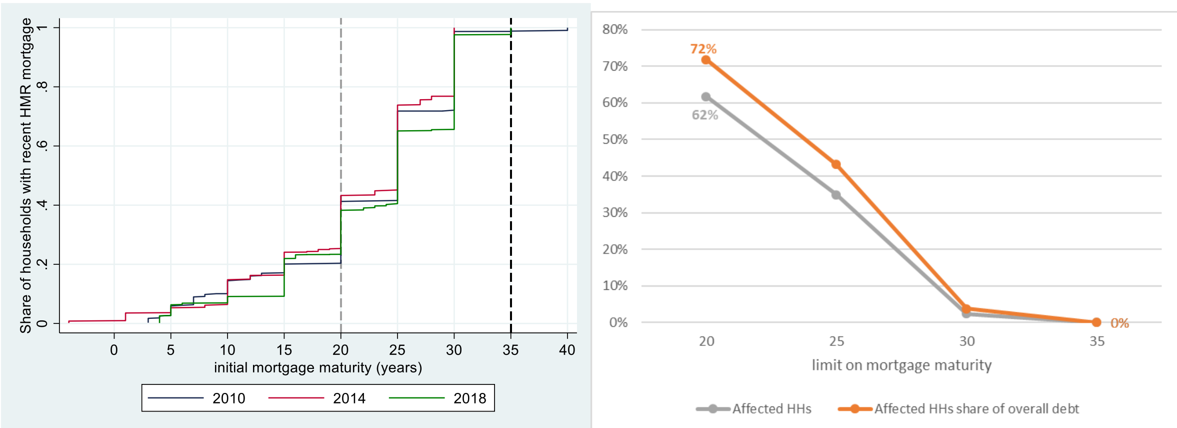

Source: Own calculations based on waves 1, 2 and 3 of the LU-HFCS; data are multiply imputed and weighted; only households with recent HMR mortgages.
Panel (a): cumulative distribution functions are calculated across all 5 implicates for each year. Vertical lines indicate lowest and highest limits envisaged by the law.

In summary, data from the 2018 LU-HFCS survey does not provide evidence of a general worsening of debt burden ratios among households resident in Luxembourg. Increases in the median LTV ratio or the median debt-to-income ratios were not statistically significant. The median debt service-to-income (DSI) ratio declined marginally, possibly reflecting the reduction in mortgage rates. The relative stability of the DSI ratio may reflect the rapid increase in the share of fixed-rate HMR mortgages[16] from 25% of the outstanding stock in 2014 to 38% in 2018 based on estimates from the LU-HFCS survey.

15 Some less regular mortgage maturities are also visible, reflecting some irregular values as reported by survey participants.
16 Fixed-rate mortgages usually involve higher interest rates.

3 Household financial vulnerability

3.1 Financial vulnerability in the 2018 baseline

The debt burden ratios in Table 2 only provide a limited measure of household financial vulnerability. Following established practice in the literature[17], we measure a household *"probability of default"* (PD) by combining household monthly net income, debt service payments, basic living expenditures and liquid asset holdings[18]. This indicator does not measure the probability of default derived from a Merton-type model or an empirical frequency in the data. However, given the absence of a credit register, it allows us to use survey data to identify financially vulnerable households, meaning those that could face serious problems servicing their monthly debt payment following an economic shock.[19]

As in Giordana and Ziegelmeyer (2020), we define household *i*'s probability of default (PD_i) as a function of its monthly financial margin (FM_i) and its liquid asset holdings[20] (LIQ_i):

$$PD_i = \begin{cases} 0 & if\ FM_i \geq 0\ \ or\ \ |FM_i| \cdot 3 \leq LIQ_i \\ 1 - \dfrac{LIQ_i}{|FM_i| \cdot 3} & if\ \ FM_i < 0\ and\ |FM_i| \cdot 3 > LIQ_i \end{cases}$$

Following the literature, the financial margin is measured as household gross income minus taxes, social security contributions, regular debt service payments, rent paid (zero for our sample of households with recent HMR mortgages) and basic living costs[21]. A household's probability of default is set to zero if it has a positive financial margin or if its liquid assets are sufficient to cover its negative financial margin for at least three months (matching the conventional 90-day limit used to define non-performing loans). Otherwise, the household's probability of default is a simple function of its financial margin and its liquid assets. Thus, this indicator focuses on liquidity risk and measures the probability that the household falls behind in its debt payments.

Also following Giordana and Ziegelmeyer (2020), the bank's loss given default associated with household i (LGD_i) is the difference between the household's total debt (D_i) and real estate assets (A_i) after applying a 25% haircut, weighted by the household probability of default:

$$LGD_i = PD_i(D_i - 0.75 \cdot A_i).$$

As household defaults are not observed in our sample, the LGD measure should be understood as an expected amount. We classify a household as financially vulnerable if its LGD exceeds zero. Note that a positive probability of default is a necessary condition for a household to be financially vulnerable on this definition. In a robustness exercise, we use the simpler criterion that financially vulnerable households are all those with a positive probability of default (see Appendix B).

---

17 See Albacete and Fessler (2010); Ampudia et al. (2016); Giordana and Ziegelmeyer (2020); Meriküll and Rõõm (2020).

18 Leika and Marchettini (2017) call this measure "probability of incurring distress".

19 See Hallissey et al. (2014) and Nier et al. (2019) for studies using loan-level data to help calibrate a macro-prudential tool. Since these authors observe defaults directly, they do not need model-generated default rates or survey-based measures.

20 We measure liquid assets as the sum of bank deposits (mainly sight and saving accounts), stocks (publicly traded stocks, mutual funds invested mostly in stocks, managed accounts, hedge funds), bonds, and potentially less liquid assets (including private businesses other than self-employment and other assets).

21 Basic living costs are estimated using household specific amounts spent on utilities (e.g., electricity, water, gas, telephone...) and on food consumed at home, as well as 50% of the amounts spent on food outside the home.

Table 4 reports the share of financially vulnerable households in 2014 and 2018 using both definitions. Among households with recent HMR mortgages, in 2018 only 2.7% had a positive PD and 1% had a positive LGD. Among households with any recent mortgages (OREP as well as HMR), 3.1% had a positive PD and 1.7% had a positive LGD. Given the limited number of households with recent mortgages, vulnerability measures are not analysed by household characteristic.

The benign economic environment in 2018 suggests a very low share of households with a PD>0 or a LGD>0. However, even in these conditions macro-prudential instruments should focus on those households who may become financially stressed if economic conditions deteriorate.

Table 4: Share of financially vulnerable households - baseline

	Wave 2014		Wave 2018		p-value of difference btw.
Vulnerability measures	Mean	Std. Err.	Mean	Std. Err.	2018 and 2014
	Households with recent HMR mortgage				
Probability of default > 0	5.6%	2.5%	2.7%	2.4%	0.45
Loss given default > 0	0.9%	†	1.0%	†	0.95
	Households with recent HMR or OREP mortgage				
Probability of default > 0	5.9%	2.3%	3.1%	2.2%	0.42
Loss given default > 0	0.6%	†	1.7%	1.3%	0.51

*Source: Own calculations based on waves 2 and 3 of the LU-HFCS; data are multiply imputed and weighted; variance estimation based on 1000 replicate weights. P-values indicate whether difference between 2014 and 2018 is significant: *** p<0.01, ** p<0.05, * p<0.1. † std. err. not available because in one implicate no household is vulnerable.*

3.2 Financial vulnerability in the adverse scenario

Building on Giordana and Ziegelmeyer (2020), we simulate household balance sheets in an adverse economic scenario to identify which households are more vulnerable to financial stress. We estimate a logit model for the probability that an individual is unemployed and use it to simulate the unemployment status of each active household member covered by the 2018 wave of the LU-HFCS (Appendix C for details). We extend Giordana and Ziegelmeyer's (2020) simulation method in two ways: i) the unemployment shock is simulated at the individual level (instead of the household level). ii) the logit model for the probability of an individual being unemployed includes additional explanatory variables, such as language skills, sector of employment and labour status.

For any individual who becomes unemployed, we set their individual labour income to zero and recalculate the household financial margin, probability of default and loss given default. Luxembourg unemployment benefits cover employees and self-employed workers for a maximum of twelve months (with certain exceptions), replacing a substantial share of lost labour income.[22] Therefore, our adverse scenario assumes that selected individuals would be unemployed for more than one year. Household income is obtained by summing across all household members, who may be affected differently by the

[22] Unemployment benefits in Luxembourg are 80% of the last gross wage or 85% if the unemployed person receives child benefits for his/her dependents. Unemployment benefits cannot exceed 2.5 times the minimum wage during the first 6 months and 2.0 times the minimum wage in the following 6 months.

unemployment shock. Simulated household net income is not allowed to fall below the *revenu d'inclusion sociale* (REVIS) (a form of guaranteed minimum income).[23]

Our adverse economic scenario assumes a 12% unemployment rate. This is an extreme shock by Luxembourg standards, more than doubling the 2019 unemployment rate. Although definitely a tail risk, this could be a plausible adverse scenario if the Covid-19 pandemic leads to an extended recession and an additional shock hits the economy (i.e Covid-19 treatments are not effective, further lockdowns, cyber-threats disrupt work-from-home arrangements) or if policy support is prematurely withdrawn.

Table 5 reports the share of households in 2018 that are financially vulnerable in the baseline and the **additional** share of households who **become** financially vulnerable in the adverse scenario. As expected, the increase is substantial in the adverse scenario, since when individuals move into unemployment, the fall in labour income reduces the household financial margin, increasing both the probability of default and the loss given default.

Table 5: Share of households that become financially vulnerable in the adverse scenario compared to the 2018 baseline

	Share of households that	
Vulnerability measures	are financially vulnerable in the 2018 baseline	become financially vulnerable in the adverse scenario
Households with recent HMR mortgage		
Probability of default > 0	2.7%	8.0%
Loss given default > 0	1.0%	3.9%
Households with recent HMR or OREP mortgage		
Probability of default > 0	3.1%	9.2%
Loss given default > 0	1.7%	4.1%

Source: Own calculations based on wave 3 of the LU-HFCS; data are multiply imputed and weighted. Reference population includes households that were financially vulnerable before the shock.

4 Evaluating different policy rules

Following the literature on household financial vulnerability,[24] we assess how effectively limits on the different debt burden ratios, or combination of ratios, can serve to identify financially vulnerable households.

[23] For a detailed description of the REVIS with respect to eligibility conditions and the amount of the benefit please consult: https://guichet.public.lu/en/citoyens/sante-social/action-sociale/aide-financiere/revenu-inclusion-sociale-revis.html#bloub-11

[24] See Albacete et al. (2018), Leika and Marchettini (2017), and Bańbuła et al. (2016).

4.1 Methodology

We implement a technique known as the "signals approach" (Kaminsky et al., 1998) or "signalling approach" (Detken et al., 2014)[25]. This implements a grid search over possible policy settings to identify those that are "optimal" in the sense that they minimise classification errors.

The confusion matrix in Table 6 illustrates the different possible outcomes when classifying households based on their debt burden ratios. The two rows of the matrix correspond to the signal from the ratio (signal triggered if above the limit, signal not triggered if below) and the two columns correspond to the condition of households (vulnerable or not vulnerable). The outcome of the classification depends on whether the signal matches the condition in the two columns. In the first column, the household is financially vulnerable, so a triggered signal results in an accurate "true positive" case in the top left cell. However, some financially vulnerable households may not trigger the signal and therefore lead to a "false negative" case in the bottom left cell. In the second column, the household is **not** financially vulnerable, so if its debt ratio triggers the signal the result is a "false positive" case in the top right cell. The bottom right cell includes households that are **not** financially vulnerable and do **not** trigger the signal, which are accurately classified as "true negative" cases.

Table 6: Confusion matrix and associated statistics

Signal	Condition		Statistics
	Vulnerable	***Not vulnerable***	
Triggered	True positive TP	False positive FP	Positive predictive value PPV = $\frac{TP}{TP+FP}$
Not triggered	False negative FN	True negative TN	Negative predictive value NPV = $\frac{TN}{FN+TN}$
Statistics	True positive rate 1-Type II error TPR = $\frac{TP}{TP+FN}$	False positive rate Type I error FPR= $\frac{FP}{FP+TN}$	Loss function = $\theta(1-TPR)+(1-\theta)FPR$, with $\theta \in (0,1)$ Markedness = $PPV+NPV-1$

Type I classification errors correspond to "false positive" cases in the second column (households incorrectly classified as vulnerable). Type II errors correspond to "false negative" cases in the first column (households incorrectly classified as not vulnerable). Varying the limit that triggers the signal will reveal a trade-off between these two types of classification errors. For instance, higher limits will reduce the number of households classified as vulnerable and therefore increase type II errors (missing some vulnerable households). Lower limits will raise the number of households classified as vulnerable and therefore increase type I errors (identifying too many households as vulnerable). The Positive Predictive Value (PPV) measures the probability that a household is financially vulnerable when the signal is triggered. Similarly, the Negative Predictive Value (NPV) measures the probability that a household is **not** financially vulnerable when the signal is **not** triggered. PPV and NPV reflect the trade-

25 This approach has also been used to evaluate indicators of economic recessions and expansions (Berge and Jordà, 2011), to evaluate the performance of investment strategies (Jordà and Taylor, 2011), to evaluate indicators of real credit contractions and expansions (Jordà, 2012), and to evaluate early-warning systems for bank distress (Betz et al. 2014), for banking crises (Drehmann and Juselius, 2014), or for financial crises (Candelon et al. 2012; Detken et al. 2014).

off between type I and type II errors. The cell in the bottom right corner defines the Loss function, which combines the True Positive Rate with the False Positive Rate, as well as the Markedness statistic, which combines the Positive Predictive Value and the Negative Predictive Value. Below, we will focus on the Loss function, although Candelon et al. (2012) provide a more detailed discussion of the relative advantages and disadvantages of these different statistics.

The signals approach relies on three elements: the binary condition variable, the criteria to evaluate the performance of different classification rules and the criteria to identify the "optimal" limit for each rule.

(a) The condition variable takes the value one for financially vulnerable households and zero for non-vulnerable households. Below, we define vulnerable households as those with a nonzero loss given default. In Appendix B we present a robustness exercise defining vulnerable households as all those with a positive probability of default.

(b) To compare classification rules, we use the Receiver Operating Characteristic curve (ROC[26]). In particular, the Area under the ROC curve (AUROC), a statistic ranging from 0 to 1, summarises the performance of each classification rule across all candidate limits.[27] A classification rule with an AUROC of 1 is perfectly informative, while one with an AUROC of 0.5 is uninformative. An AUROC significantly lower than 0.5 may be informative if it indicates a systematic inverse relationship between the signal and the underlying condition. Performance within a subset of candidate limits can be evaluated with the "partial" AUROC, normalised so that it also ranges from 0 to 1.

(c) Criteria to identify "optimal" limits

    Following Detken et al (2014), we minimise the loss function (bottom right cell in Table 6) to find the "optimal" limit[28]. Varying the loss function parameter θ allows one to consider different policy preferences regarding the trade-off between type I and type II errors.

4.2 Evaluation results

In this section, we apply the signals approach to evaluate how well limits on individual debt burden ratios or combinations of these ratios[29] identify financially vulnerable households. We perform a grid search evaluating the loss function over a set of possible values of the limit (including the legal range). For each ratio, the grid ranges from around the 5th percentile to around the 95th percentile of the distribution observed in the data. Thus, for the LTV ratio, we consider 37 values in the interval [0.05; 1.9]. For the MDI ratio and the DI ratio, we consider 80 values in the interval [0.5; 20.5]. For the DSI ratio, we consider 100 values in the range [0.1; 1.1]. This number of points ensures sufficient granularity to identify the "optimal" limit with some precision. At each of these values, we calculate

[26] The Receiver Operating Characteristic (ROC) curve plots the True Positive Rate (TPR) against type I error for all candidate limits in our grid search. See Hanley and McNeil (1982) for a definition.

[27] See Hsieh and Turnbull (1996) for further definitions.

[28] This limit will only be "optimal" in the sense of minimising classification errors. In our application, this will not necessarily correspond to minimising the policymaker's loss function, which would require specifying the links between classification errors, systemic risk and social welfare.

[29] Mortgage maturity is not really a debt burden ratio (section 2.2), so we only analyse it in combination with other ratios.

the number of type I errors and type II errors to evaluate three alternative loss functions (based on different values of the weight θ). This allows us to identify the "optimal" value of the limit as the one that delivers the lowest value for the given loss function. This "optimal" limit may vary across different loss functions, but unfortunately, there is no obvious best choice of the weight θ, which should reflect policy makers' preferences regarding type I and type II errors.

4.2.1 Baseline

Table 7 reports the results under the 2018 baseline. In the first column, the Area Under the Receiver Operating Characteristic curve (AUROC) assesses each ratio's performance over the entire range of candidate limits considered.[30] The AUROCs from the DI and MDI ratios are not statistically different, so we cannot reject the hypothesis that they are equally informative[31]. The DSI and LTV ratios perform significantly worse but their AUROCs are still well above 0.5, suggesting that even these ratios are informative about household financial vulnerability. The second column focuses only on the range of possible values envisaged by the law, reporting partial AUROCs (also normalised to range from 0 to 1). These are all higher than the AUROCs, suggesting that the ratios perform better in the legal range than over the wider range considered in the first column.

Table 7: Performance of individual debt burden ratios in the 2018 baseline

Ratio	AUROC (s.d.)[a]	Partial AUROC[b] (s.d.)[a]	Weight on Type II error (θ)	Optimal limit[c]	FPR	1-TPR	Loss	PPV	NPV	Marked.
LTV	0.720 (0.0099)	0.727 (0.0098)	0.25	1.85	0.041	0.679	0.200	0.072	0.993	0.065
			0.5	1.1	0.165	0.390	0.278	0.035	0.995	0.031
			0.75	0.8	0.613	0.000	0.153	0.016	1.000	0.016
MDI	0.924 (0.0062)	0.985 (0.0029)	0.25	14.75	0.050	0.101	0.063	0.151	0.999	0.150
			0.5	14.75	0.050	0.101	0.076	0.151	0.999	0.150
			0.75	9	0.206	0.000	0.052	0.046	1.000	0.046
DI	0.925 (0.0062)	0.987 (0.0027)	0.25	15	0.055	0.101	0.066	0.140	0.999	0.139
			0.5	15	0.055	0.101	0.078	0.140	0.999	0.139
			0.75	9.25	0.213	0.000	0.053	0.045	1.000	0.045
DSI	0.801 (0.0090)	0.845 (0.0083)	0.25	0.85	0.057	0.422	0.149	0.091	0.996	0.087
			0.5	0.54	0.209	0.161	0.185	0.038	0.998	0.036
			0.75	0.34	0.508	0.000	0.127	0.019	1.000	0.019

Source: Own calculations based on the 3rd wave of the LU-HFCS; data are multiply imputed and weighted; statistics are calculated by pooling all implicates. Notes: See Table 2 for definition of LTV, MDI, DI and DSI ratios and Table 6 for definition of FPR, TPR, Loss, PPV, NPV and Markedness. [a] Based on Hanley and McNeil (1982). [b] Area under the ROC curve within the range envisaged by the law for each debt burden ratio. [c] Maximum legal limits are 1 for the LTV ratio, 12 for the MDI ratio and DI ratio and 0.75 for the DSI ratio (see Table 2).

Column 3 in Table 7 indicates the three alternative values considered for the weight θ in the loss function defined in Table 6. Different weights do not always lead to different "optimal" values of the limit minimising the loss function: for DI and MDI, θ = 0.5 leads to the same "optimal" value of the limit

[30] Figure 9 in Appendix B plots the ROC curves.

[31] We construct confidence intervals using Hanley and McNeil (1982) estimates of the standard error.

as does θ=0.25. However, for LTV and DSI the "optimal" value of the limit differs depending on the weight θ, reflecting less marked changes along the ROC curve.[32] Although most ratios and weights generate few classification errors, the Positive Predicted Value is always smaller than the Negative Predicted Value, which is close to one. This reflects the fact that in the 2018 baseline only 1.0% of households are financially vulnerable (have a nonzero LGD). In most rows of Table 7, the "optimal" value of the limit is above the legal range, except when θ = 0.75 (in which case, the "optimal" limit on the DSI ratio actually falls below the legal range). However, when θ = 0.75 markedness is generally lower (last column in Table 7) suggesting that the classification is less precise. For instance, for the LTV ratio the "optimal" limit is 80% when the weight is θ = 0.75, but markedness is very low because of many false positives.

To capture the multi-dimensional aspect of household financial vulnerability, below we consider more general classification rules that combine information from several debt burden ratios simultaneously. Using a combination of ratios could change the classification errors and further reduce the loss function at the "optimal" value of the limit. We consider a classification rule that combines MDI, DSI and LTV ratios as well as mortgage maturity (MM), which was not considered above. The combined rule may trigger a signal, classifying a household as financially vulnerable, if the household exceeds the limit on any one of the four ratios considered. Alternatively, the combined rule may only trigger the signal if the household exceeds the limit on at least two of the four ratios. A less restrictive combined rule may require breaches of the limits on at least three of the four ratios. The least restrictive rule will require breaches on all four ratios. As previously, we implement a grid search, considering the same set of points as above.[33] However, since we now consider limits on several ratios simultaneously, we must explore more than 1.7 million possible combinations of policy settings.

Table 8 focuses on the combined rule requiring households to breach the limit on **at least two** of the four ratios, as this combined rule produced the highest value of the AUROC[34]. The AUROC from the combined rule in Table 8 is significantly higher than AUROCs obtained by simpler rules using DSI and LTV ratios alone (Table 7), although it is not statistically different from AUROCs based only on the MDI or DI ratios[35].

---

[32] Given our non-parametric approach, the limited number of observations and their heterogenous nature produces ROC curves that are kinked rather than smooth parabolas (see Figure 9b). As a result, outlying points on the ROC curve may be optimal for several values of the policy parameter, corresponding to different slopes of the tradeoff between Type I and Type II errors.

[33] For the MM limit, we consider 7 values over the range [5,35].

[34] Table 16 in the appendix reports full results for all combined classification rules. This also includes a robustness check using a less restrictive condition variable, when financially vulnerable households are defined as all those with a nonzero probability of default. Table 17 provides an additional robustness check by replacing MDI by DI in the set of ratios being combined.

[35] Detken et al. (2014) note that signalling performance tends to improve when combining several indicators.

Table 8: Performance of combined debt burden ratios in the 2018 baseline

Limits exceeded	AUROC (s.d.)[a]	Weight on type II error (θ)	Optimal limits*				FPR	1-TPR	Loss	PPV	NPV	Marked.
			LTV	MDI	DSI	MM						
2	0.918 (0.006)	0.25	1.90	14.93	1.08	35	0.10	0.04	0.055	0.19	1.00	0.19
		0.5	0.92	13.41	0.62	35	0.00	0.11	0.055	0.08	1.00	0.08
		0.75	0.92	13.41	0.62	35	0.00	0.11	0.028	0.08	1.00	0.08

*Source: Own calculations based on the 3rd wave of the LU-HFCS; data are multiply imputed and weighted; statistics are calculated by pooling all implicates. Notes: See Table 2 for definition of LTV, MDI, DI and DSI ratios and Table 6 for definition of FPR, TPR, Loss, PPV, NPV and Markedness. [a] Based on Hanley and McNeil (1982). * Maximum legal limits are 1 for the LTV ratio, 12 for the DI ratio and the MDI ratio, 0.75 for the DSI ratio and 35 years for MM (see Table 2).*

In addition to the higher AUROC, Table 8 also reports a lower value of the loss function and an improvement in markedness compared to using individual debt burden ratios (Table 7). Each row refers to a different weight θ, but the loss is always lower than when using individual debt burden ratios separately. Markedness only improves for $\theta = 0.25$ and $\theta = 0.5$, although the "optimal" limits fall outside the legal range (except for mortgage maturity). For $\theta = 0.75$, only the "optimal" limit on the MDI ratio falls outside the legal range.

Based on the data in the 2018 baseline, the "optimal" limits that minimise classification errors tend to be high relative to the ranges envisaged by the law. In the next sub-section, we evaluate performance when focussing on households that were not financially vulnerable in the 2018 baseline but become vulnerable in our adverse scenario.

4.2.2 Adverse scenario

The classification rules still use the individual household debt burden ratios observed in the baseline to determine the signal. However, the binary condition variable now takes the value unity only for households who were not vulnerable in the baseline but become financially vulnerable in the adverse scenario (see Table 5). Table 9 evaluates classification rules based on individual ratios and Table 10 evaluates classification rules combining different ratios.

In Table 9, the differences between AUROCs on DI, MDI and DSI ratios are not statistically significant, so performance may be comparable. While the AUROC on the LTV ratio is lower, it is statistically higher than 0.5, confirming that the LTV ratio is informative to identify financial vulnerability.

While there are more vulnerable households in the adverse scenario than the baseline, the AUROCs are lower (Table 7). This reflects a different set of vulnerable households, as well as the prospective nature of our exercise, which focuses on the weaker link between households' balance sheet condition at the time of mortgage take out and its financial vulnerability after a severe income shock. As expected, the "optimal" limits are lower than in Table 7, much closer to the legal range (see Table 2). The drawback is that type I errors are more common, reducing the AUROC and markedness. When the weight on type II error is 0.25, the "optimal" limit on the LTV ratio is the only one that exceeds the legal range. When the weight on type II error is 0.75, the "optimal" limit for the DSI ratio falls below the legal range.

Table 9: Performance of individual debt burden ratios in the adverse scenario

Ratio	AUROC (s.d.)[a]	Partial AUROC[b] (s.d.)[a]	Weight on Type II error (θ)	Optimal limit[c]	FPR	1-TPR	Loss	PPV	NPV	Marked.
LTV	0.622 (0.0052)	0.692 (0.0050)	0.25	1.75	0.048	0.948	0.273	0.043	0.961	0.004
			0.5	0.75	0.628	0.000	0.314	0.060	1.000	0.060
			0.75	0.75	0.628	0.000	0.157	0.060	1.000	0.060
MDI	0.688 (0.0051)	0.732 (0.0050)	0.25	10.25	0.133	0.565	0.241	0.116	0.974	0.091
			0.5	6.5	0.469	0.154	0.312	0.068	0.988	0.056
			0.75	5	0.673	0.000	0.168	0.057	1.000	0.057
DI	0.699 (0.0051)	0.753 (0.0049)	0.25	10.75	0.129	0.565	0.238	0.120	0.974	0.095
			0.5	7	0.417	0.154	0.285	0.076	0.989	0.065
			0.75	5	0.691	0.000	0.173	0.055	1.000	0.055
DSI	0.698 (0.0051)	0.747 (0.0049)	0.25	0.5	0.215	0.228	0.218	0.127	0.988	0.115
			0.5	0.5	0.215	0.228	0.221	0.127	0.988	0.115
			0.75	0.26	0.698	0.000	0.175	0.055	1.000	0.055

Source: Own calculations based on the 3rd wave of the LU-HFCS; data are multiply imputed and weighted; the statistics are calculated by pooling all implicates. Notes: See Table 2 for definition of LTV, MDI, DI and DSI ratios and Table 6 for definition of FPR, TPR, Loss, PPV, NPV and Markedness. [a] Based on Hanley and McNeil (1982). [b] Area under the ROC curve within the legal range of the debt burden ratio. [c] Maximum legal limits are 1 for the LTV ratio, 12 for the DI/MDI ratio and 0.75 for the DSI ratio (see Table 2).

Table 10 focuses on the combined classification rule that requires at least three limits to be breached, since this rule yields the highest AUROC in the adverse scenario[36] and includes the only case with all "optimal" limits within their legal range (θ=0.75). Markedness reaches 0.43 in this row, which is remarkably high compared to classification rules based on individual ratios (Table 9), as well as compared to the 2018 baseline (Table 7 and Table 8). In the row marked θ=0.5, the "optimal" limit for the MDI ratio is above the legal range and the one for the LTV ratio is below the legal range. In the row marked θ=0.25, the MM limit is 17 years, also below the legal range. Of course, in choosing the ranges envisaged by the law, legislators may not have considered limiting several ratios simultaneously. In addition, it seems unlikely that they agreed on a common loss function with explicit weights on type I and type II errors.

Table 10: Performance of combined debt burden ratios in the adverse scenario

Rule	AUROC (s.d.)[a]	Weight on Type II error (θ)	Limits* LTV	MDI	DSI	MM	FPR	1-TPR	Loss	PPV	NPV	Marked.
3	0.705 (0.005)	0.25	0.87	5.06	0.48	17	0.00	0.42	0.105	0.09	1.00	0.09
		0.5	0.62	20.50	0.49	23	0.24	0.13	0.181	0.20	0.99	0.19
		0.75	0.98	7.34	0.57	35	0.52	0.02	0.149	0.45	0.98	0.43

*Source: Own calculations based on the 3rd wave of the LU-HFCS; data are multiply imputed and weighted; the statistics are calculated by pooling all implicates. Notes: See Table 2 for definition of LTV, MDI, DI and DSI ratios and Table 6 for definition of FPR, TPR, Loss, PPV, NPV and Markedness. [a] Based on Hanley and McNeil (1982). * Maximum legal limits are 1 for the LTV ratio, 12 for the DI/MDI ratio, 0.75 for the DSI ratio and 35 years for MM (see Table 2).*

[36] Table 18 in the appendix evaluates the performance of all combined rules in the adverse scenario. Table 19 provides a robustness check by replacing DI with MDI in the set of ratios being combined.

The high AUROCs confirm that debt burden ratios observed in the benign baseline conditions are informative to identify households that only become financially vulnerable following the adverse shock. As in the baseline, rules that combine several ratios outperform rules based on individual ratios.

Overall, the “optimal” limits depend on the policymaker objective as represented by the weights on different types of classification errors and on the built-in “stress” in the adverse scenario. Our results indicate that the legal range is sufficiently large to accomodate policy calibrations aimed at enhancing household and bank resilience to a vast array of negative shocks.

5 Allowing for household heterogeneity

In the previous section, the signals approach assigned the same weight to every household, regardless of its individual situation. This implicitly assumed that the cost of misclassification is the same across households, which tends to weaken the link between minimising classification errors and the policymaker objective of maximising social welfare. In this section, we recognise that individual households may contribute differently to social costs and benefits, depending on their individual volume of debt, probability of default or loss given default. First, we compare average loss given default for households below and above the limit at different policy settings. Second, we plot cost frontiers comparing the share of mortgage debt “incorrectly” granted to the share of debt “incorrectly” denied at different policy settings. Our paper cannot establish whether net effects on social welfare are positive or negative, because this requires comparing the net present value of one euro of debt incorrectly denied to the net present value of one euro of debt incorrectly granted. Since our framework is static, we can only compare the volume of debt incorrectly denied to the volume of debt incorrectly granted. This can also be expressed as the LGD of households who comply with the limit and the LGD of households who would be rationed by the limit. These elements are the starting point for the net present value comparison, but are not sufficient to measure social welfare.

To account for household heterogeneity, one could select “optimal” limits on the debt ratios by comparing average LGD across affected households and across compliant households at different policy settings. In principle, the average LGD across affected households should be higher at higher limits, while the average LGD across compliant households should remain close to zero. Type I error (erroneously classifying households as vulnerable) reduces the average LGD across affected households (as they will include more non-vulnerable households) and type II error (erroneously classifying households as not vulnerable) raises the average LGD across compliant households (as they will include more vulnerable households). Therefore, both errors reduce the gap across groups in terms of average LGD and therefore, the peaks in Figure 6 and Figure 7 indicate optimal limits for this criterion. As a result, it is not surprising that limits on the debt-to-income ratios, which minimized classification errors (largest AUROC), also feature the largest gaps between affected and compliant households in terms of average LGD (panels (b) and (c) in Figure 6 and Figure 7).

In the 2018 baseline (Figure 6), the gap is systematically largest at the top of the legal range (most “optimal” limits in Table 7 were even higher). However, for the DSI ratio in panel (d) the gap reaches a local peak within the legal range at 54%, which was the “optimal” limit in Table 7. This local peak reflects differences in LGD across households as well as type II classification errors, since a handful of vulnerable households with relatively high LGDs have MDI ratios below the 55% limit and therefore switch to the compliant group at higher limits. However, the gap in average LGD still reaches an overall maximum at 75%, the top of the legal range.

Figure 6: Average loss given default of compliant and non-compliant households with recent HMR mortgages in 2018

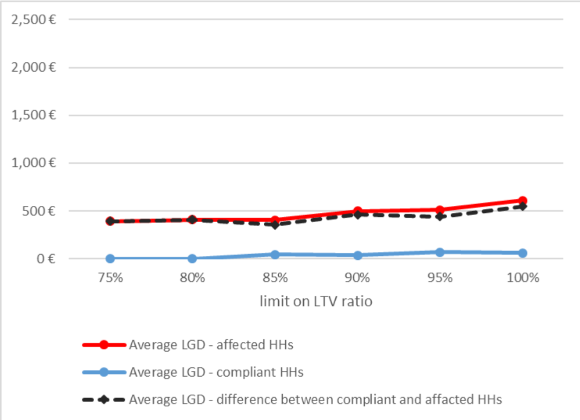

(a) LTV ratio

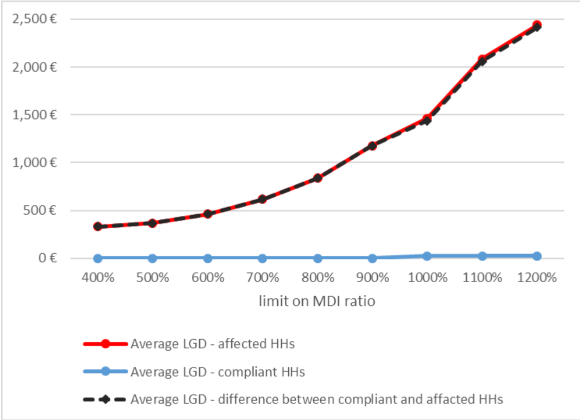

(b) MDI ratio

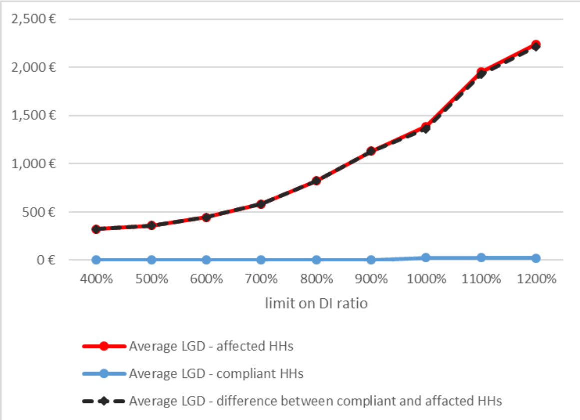

(c) DI ratio

(d) DSI ratio

Source: Own calculations based on the 3rd wave of the LU-HFCS; data are multiply imputed and weighted; the statistics are calculated by pooling all implicates. Note: Peaks in average LGD difference between compliant and affected households (dashed black line) indicate best limits.

Figure 7 reports the gap in average LGDs in the adverse scenario, which is generally lower than it was in the 2018 baseline (Figure 6), reflecting the poorer performance identifying vulnerable households (lower AUROCs in Table 9). In the adverse scenario, the gap in average LGDs peaks within the legal range for MDI, DI and DSI, although above the "optimal" limits identified by minimising classification errors (Table 9). At the right end of the charts, average LGDs increase after a local minimum for MDI, DI and DSI, which reflects type I classification errors at the right tail (the blue lines are steady and red lines increase). For the LTV ratio, the gap is largest at the top of the legal range (100%), which is substantially above the "optimal" limit from the signals approach (75%). However, for the LTV ratio changes in the value of the limit have little impact on the gap in average LGDs.

In the baseline scenario, household heterogeneity in terms of LGD does not change the conclusions of the signals approach. However, in the adverse scenario, accounting for LGD heterogeneity across households suggests "optimal" limits could be higher as additional classification errors at higher levels of the limit have a smaller impact on the gap in average LGDs.

Figure 7: Average loss given default of compliant and non-compliant households with recent HMR mortgages in the adverse scenario

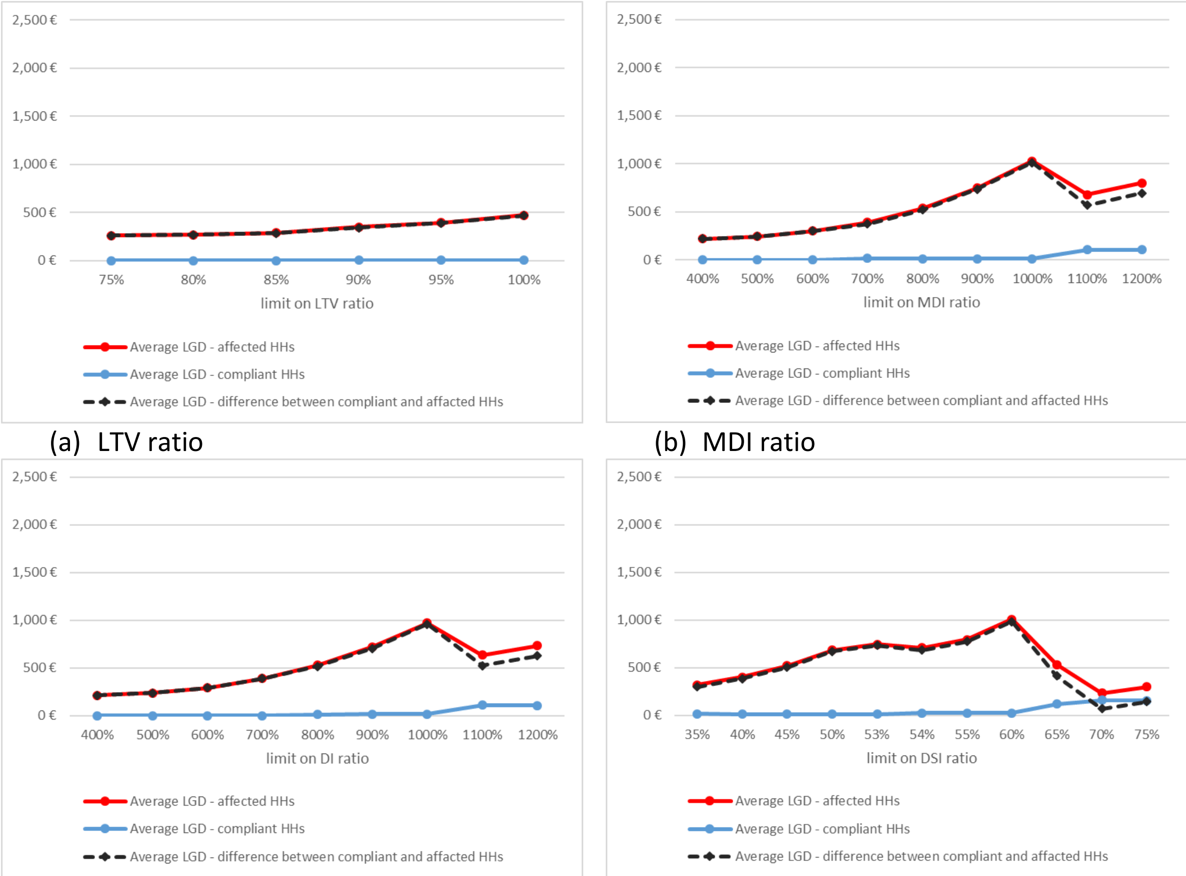

(a) LTV ratio

(b) MDI ratio

(c) DI ratio

(d) DSI ratio

Source: Own calculations based on the 3rd wave of the LU-HFCS; data are multiply imputed and weighted; the statistics are calculated by pooling all implicates. Note: Peaks in average LGD difference between compliant and affected households (dashed black line) indicate best limits.

Next, we plot a cost frontier by weighting each incorrectly classified household by the amount of its most recent mortgage. The cost frontier plots the share of total mortgage volume that would be incorrectly denied under the policy (type I error) against the share of total mortgage volume that would be granted incorrectly (type II error).[37] The closer the cost frontier is to the origin, the smaller the costs of classification errors. In principle, type I errors could be transformed into the (net present value of) economic costs from mortgages incorrectly denied by taking into account multiplier effects of economic activity in the construction industry. Likewise, type II errors could be transformed into the (net present value of) costs of additional systemic risk from mortgages incorrectly granted. However, if we are willing to accept the plausible assumption that these transformations are monotonic, then

[37] We adapted Fawcett's (2006) proposal to focus on the costs of classification errors instead of plotting benefits of true positives against cost of false negatives.

we can still interpret these cost frontiers without actually needing to choose an explicit specification for each transformation.

Figure 8 depicts the cost frontiers from the 2018 baseline and adverse scenario. For the individual ratios, the markers indicate the outcomes of selected limits within the legal range. For the policy rule combining different ratios, the markers show the outcome for different values of θ. To compare points across the cost frontiers, we calculate the distance of each point from the origin. The smaller the distance the better the performance, assuming the social cost of one euro of mortgage incorrectly granted is equivalent to the social cost of one euro of mortgage incorrectly denied. Note that visual comparisons can be misleading, since the scale of the x-axis is larger than the scale of the y-axis, particularly in Figure 8(a).

Figure 8: Cost frontiers households with recent HMR mortgages

(a) 2018 Baseline

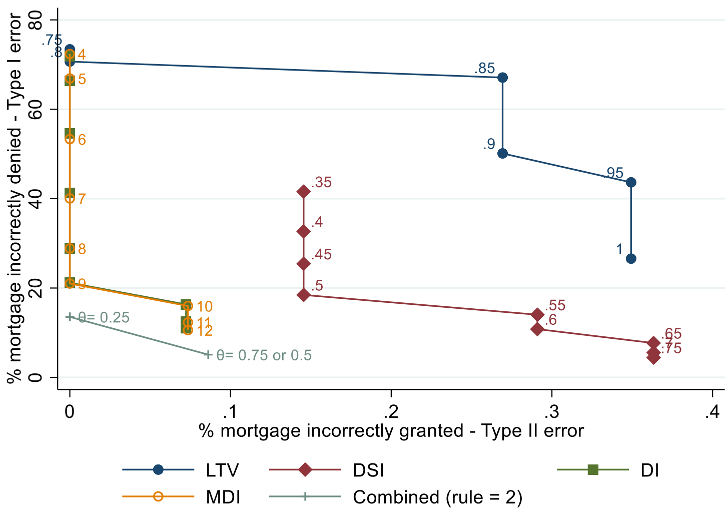

(b) Adverse scenario

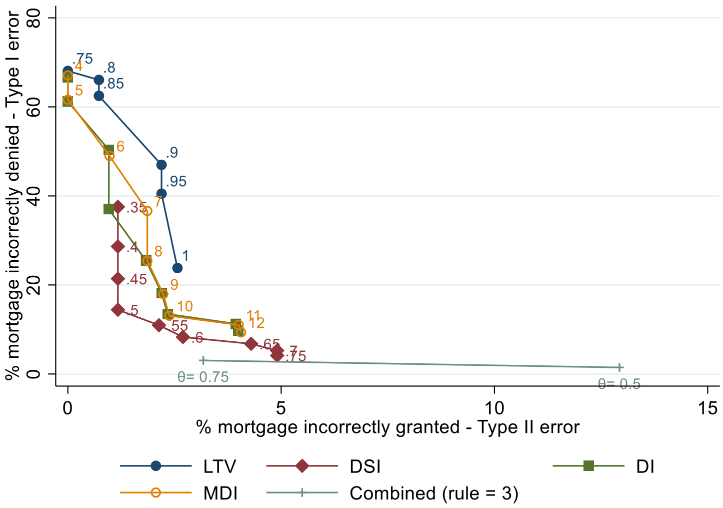

Source: Own calculations based on the 3rd wave of the LU-HFCS; data are multiply imputed and weighted; classification errors are calculated by pooling all implicates.

In the 2018 baseline, the cost frontiers tend to reflect the ranking by AUROCs in Table 7, with the MDI ratio and the DI ratio (largely overlapping in Figure 8) performing better than the DSI ratio or the LTV ratio. However, at several points the cost frontier for the DSI ratio is closer to the origin than the MDI and DI frontiers, indicating lower costs. In particular, limits above 50% on the DSI ratio outperform limits on the other ratios. For the 75% limit on the DSI ratio, the cost frontier is closer to the origin than the policy combining different ratios (bottom row of Table 8, θ=0.75).

In the adverse scenario, limits on the DSI ratio clearly perform better than limits on the other individual ratios, as the associated frontier is closer to the origin. However, the policy combining ratios performs even better (with θ=0.75, see "optimal" limits in Table 10).

To allow for the case where the social cost of mortgages incorrectly denied differs from the social cost of mortgages incorrectly granted, we define a new loss function as a linear combination of the percentage of mortgages incorrectly denied and the percentage of mortgages incorrectly granted. Table 11 reports the resulting "optimal" limits that minimise this loss function for three alternative weights θ on the share of mortgages incorrectly granted.

In the 2018 baseline, the "optimal" limits are at the top of the legal range for every policy based on an individual ratio, regardless of the weight chosen. In the adverse scenario, when θ equals 0.25 or 0.5 the "optimal" limit for the DSI ratio is 75%, but when θ = 0.75 the "optimal" limit is 60%. For the MDI ratio and the DI ratio the optimal limit is 1200% when θ equals 0.25 or 0.5, and 1000% when θ = 0.75. These "optimal" limits, which minimize the linear combination of the percentage of mortgages incorrectly denied and the percentage of mortgages incorrectly granted, differ substantially from those identified in Table 9, which did not account for heterogeneity at the household level. However, they are similar to the "optimal" limits identified in Figure 7 by maximising the gap in average LGDs.

Table 11: Limits on individual ratios that minimize a linear combination of the volume of mortgages incorrectly denied and mortgages incorrectly granted

Weight on Type II error (θ)	2018 Baseline				Adverse scenario			
	LTV	DSI	MDI	DI	LTV	DSI	MDI	DI
0.25	1.00 (23.9)	0.75 (4.5)	12 (10.6)	12 (11)	1.00 (23.9)	0.75 (6.4)	12 (10.2)	12 (10.6)
0.5	1.00 (23.9)	0.75 (4.5)	12 (10.6)	12 (11)	1.00 (23.9)	0.75 (6.4)	12 (10.2)	12 (10.6)
0.75	1.00 (23.9)	0.75 (4.5)	12 (10.6)	12 (11)	1.00 (23.9)	0.6 (8.7)	10 (13.3)	10 (13.7)

Source: Own calculations based on the 3rd wave of the LU-HFCS; data are multiply imputed and weighted; the statistics are calculated by pooling all implicates.
Note: The distance from the origin in Figure 8 appears in parenthesis.

6 Discussion

In this section, we discuss how our analysis can support the use of borrower-based measures in Luxembourg. The policy evaluation above depends on the chosen definition of household financial vulnerability, as well as policymaker preferences with respect to type I and type II classification errors. Since the approach relies on explicit assumptions, it can provide a useful basis to compare alternative instrument settings. However, it does require policymakers to agree on a definition of household financial vulnerability and decide how to weight type I and type II classification errors. Modifying the definition or the relative weight on type I and type II errors can lead to different "optimal" settings of the borrower-based measures.

Our analysis finds that the LTV ratio does not always perform best in identifying financially vulnerable households. However, our analysis abstracts from many practical considerations. For example, some data is directly observable by the lender (i.e. mortgage and property value used to calculate the LTV ratio) while other information is less accessible (e.g. debt outstanding, debt service on older debt) or needs to be estimated (e.g. disposable income). In practice, borrower-based instruments should complement lenders' internal credit-scoring evaluation of individual borrowers, rather than replace this practice. These internal evaluations are likely to consider debt burden ratios similar to those in the legal text. This provides another reason for authorities to focus on the LTV ratio, which could be consistently calculated by every lender.

Our results support the recommendation that the regulatory limit on the LTV ratio be considered a complement to lenders' internal credit-scoring practices. First, when focussing on households who only become financially vulnerable after an adverse income shock, we find that borrower-based measures are best targeted when a limit on the LTV ratio is combined with limits on other debt burden ratios. Second, the "optimal" limit for distinguishing vulnerable and non-vulnerable households depends on economic conditions as well the loss function chosen. In particular, our analysis of the 2018 baseline would suggest that the "optimal" limits can be fairly high, constraining credit to fewer households. Instead, in the adverse scenario the "optimal" limits are lower, constraining credit to more households. When we consider combining limits on different ratios, the "optimal" limit on the LTV ratio is 98%, very close to regulatory limit currently imposed on first-time buyers[38].

7 Conclusion

This paper proposes a data-driven approach to borrower-based macro-prudential policy using data on individual households. Thus, it can serve as a complement to analyses based on macro-economic models or bank-level data. First, we investigate household financial vulnerability using balance sheet information from the Luxembourg Household Finance and Consumption Survey. Between 2014 and 2018, the median loan-to-value ratio increased 9 percentage points and the median mortgage debt-to-income ratio increased more than 60 percentage points, but neither change was statistically significant. The share of financially vulnerable households also did not change significantly. In the 2018 data, only 1.0% of households with a recent mortgage on their main residence had a nonzero loss given default and only 2.7% had a nonzero probability of default (based on their financial margin). Simulating an adverse scenario, the share of vulnerable households increases substantially: the share with a nonzero loss given default increases by 3.9 percentage points and the share with a nonzero probability of default increases by 8 percentage points.

[38] Our sample focuses on households who bought their main residence between 2015 and 2018. Since the reference person is less than 45 years old in 70% of these households, they are probably first-time buyers, for whom the LTV limit is set at 100%.

To evaluate the ex ante impact of borrower-based measures, we consider imposing limits on several household debt ratios, as envisaged by the law. For households with recent HMR mortgages in 2018, we calculate the share whose debt burden ratios would exceed the limit at different policy settings, as well as the share of overall debt in this group that is held by affected households and the reduction in their overall debt required to comply with the limit. Focusing on the LTV ratio, the least restrictive limit envisaged by the law (100%) would have affected 24% of households who borrowed to buy their main residence between 2015 and 2018. Affected households held 28% of overall debt in this group, and would have had to reduce their debt by 7% to comply with the limit. Had the most restrictive LTV limit envisaged by the law (75%) been applied, credit would have been rationed for 64% of households with recent mortgages. These affected households represented 74% of overall debt in this group and would have had to reduce their debt by 18% to comply with the limit.

Following Albacete et al. (2018) and Bańbula et al. (2016), we implement the "signals approach" to evaluate how debt burden ratios, individually or in combination, can serve to identify financially vulnerable households. This approach can serve to calculate the limits on debt burden ratios that are "optimal" in the sense of minimising classification errors, which. However, results depend on the definition of financial vulnerability and the relative weight on different error types. In general, we show that the loan-to-value ratio can improve the identification of financially vulnerable households when used in combination with other debt burden ratios.

Based on our results, the "optimal" limits (in terms of identifying vulnerable households) are often above the range envisaged by the law, reflecting the favourable economic conditions when the most recent survey wave was collected in 2018. However, macro-prudential policy should aim to build resilience to adverse shocks. Therefore, we simulate an adverse scenario and find that many of the "optimal" limits then fall in the range envisaged by the law.

Two caveats of the signals approach are particularly relevant from the policy perspective. First, this approach does not account for several practical considerations. Measuring household disposable income is likely to be particularly challenging, since financial situations differ dramatically across households and lenders may define disposable income differently, opening the way to controversy. Instead, the LTV ratio focuses on the collateral posted in a financial transaction, considering the individual asset and liability in isolation from the household financial situation. Therefore, the definition of the LTV ratio is less controversial, facilitating comparisons across households and lenders.

Second, the signals approach focuses on classifying households into two groups (vulnerable and non-vulnerable) and assumes that classification errors always have the same effects, regardless of household characteristics. Therefore, we evaluate the impact of different policy settings on households' loss given default or the volume of household debt. Accounting in this way for the heterogeneous distribution of household debt, collateral and income can lead to quite different "optimal" limits.

Since Luxembourg does not yet have a credit register, the Household Finance and Consumption Survey provides an indispensable source of household-level data for ex-ante policy evaluation. However, the limited number of observations means that point estimates are subject to significant uncertainty. Therefore, our results should be interpreted carefully and considered as complementary to other policy analyses.

To conclude, our analysis finds that the "optimal" setting of borrower-based instruments will vary with the definition of household financial vulnerability, with policymaker preferences with respect to the trade-off between type I and type II errors, and with underlying economic conditions. Our analysis provides a framework to assist in policy design by rendering explicit the definitions, assumptions and scenarios that need to be adjusted to reflect policymaker preferences or evolving economic conditions.

In future work, one could extend the approach in three directions. First, one can build on the literature on machine learning (Fawcett (2006); Kim et al. (2012); Garrido et al. (2018)) to better account for heterogeneity across households. Second, one can analyse the short-run effects of borrower-based measures on the price of housing in alternative partial equilibrium scenarios (Makdissi and Wodon (2002); Giordana (2016)). Finally, one can build on Gross and Población Garcia (2017) and combine the signals approach with a dynamic micro-macro stress test model.

8 References

Albacete, N., and P. Fessler (2010): Stress Testing Austrian Households. Österreichische Nationalbank Financial Stability Report 19, June.

Albacete, N., Fessler, P., and P. Lindner (2018): One policy to rule them all? On the effectiveness of LTV, DTI and DSTI ratio limits as macroprudential policy tools. Financial Stability Report 35, Österreichische Nationalbank, June.

Albacete, N., and P. Lindner (2013): Household Vulnerability in Austria – A Microeconomic Analysis Based on the Household Finance and Consumption Survey. Financial Stability Report 25, Österreichische Nationalbank, June.

Albacete, N., and P. Lindner (2017): Simulating the impact of borrower-based macroprudential policies on mortgages and the real estate sector in Austria - evidence from the Household Finance and Consumption Survey 2014. Financial Stability Report 33, Österreichische Nationalbank, June.

Ampudia, M., H. van Vlokhoven, and D. Zochowski (2016): Financial fragility of euro area households. Journal of Financial Stability, 27 , 250–262.

Bańbuła, P., Kotuła, A., Przeworska J., and P. Strzelecki (2016): Which households are really financially distressed: How micro data could inform the macroprudential policy. IFC Bulletins chapters, Bank for International Settlements (eds.), Combining micro and macro statistical data for financial stability analysis, vol. 41, Basle.

Berge, T.J., and Jordà, Ò. (2011): Evaluating the classification of economic activity into recessions and expansions, American Economic Journal: Macroeconomics, 3(2), 246-277.

Betz, F., Oprică, S., Peltonen, T.A., and P. Sarlin (2014): Predicting distress in European banks. Journal of Banking and Finance, 45, 225-241.

Candelon, B., Dumitrescu, E., and C. Hurlin (2012): How to evaluate an early-warning system: Toward a unified statistical framework for assessing financial crises forecasting methods. IMF Economic Review, 60(1), 75-113.

CGFS (2016): Experiences with the ex ante appraisal of macroprudential instruments. CGFS Paper No. 56.

Chen, J. Columba, F., and C. Beaumont (2016): Macroprudential and monetary policy interactions in a DSGE model for Sweden, IMF Working Paper No. 16/74.

Chen, Y., Guarda, P., Mathä, T., Pulina, G., and M. Ziegelmeyer (2021): BCL Bulletin 2021-1, Analyse 1, 77-84.

Chen, Y., Mathä, T., Pulina, G., Schuster, B., and M. Ziegelmeyer (2020): The Luxembourg Household Finance and Consumption Survey: Results from the third wave. Banque centrale du Luxembourg Working Paper 142.

Detken, C., Weeken, O., Alessi, L., Bonfim, D., Boucinha, M., Castro, C., Frontczak, S., Giordana, G., Giese, J., Jahn, N., Kakes, J., Klaus, B., Lang, J., Puzanova, N., and P. Welz (2014): Operationalising the countercyclical capital buffer: indicator selection, threshold identification and calibration options. ESRB Occasional Paper Series, 5, 1-95.

Drehmann, M., and M. Juselius (2014): Evaluating early warning indicators of banking crises: Satisfying policy requirements. International Journal of Forecasting, 22(3), 493-518.

Fawcett, T., (2006): ROC graphs with instance-varying costs. Pattern Recognition Letters, 27(8), 882-891.

Ferreira Filipe, S. (2018): Housing prices and mortgage credit in Luxembourg. BCL Working Paper 117. February 2018.

Garrido, F., Verbeke, W, and C. Bravo (2018): A Robust profit measure for binary classification model evaluation. Expert Systems with Applications, 92, 154-160.

Giordana, G. (2016): Welfare and stochastic dominance for the measurement of banks' domestic systemic importance: analytical framework and application. International Journal of Finance and Economics, 21, 192–208.

Giordana, G., and M. Ziegelmeyer (2017): Household debt burden and financial vulnerability in Luxembourg. BCL Working Paper 113. Published in IFC Bulletins chapters, Bank for International Settlements (ed.), Data Needs and Statistics Compilation for Macroprudential Analysis, 2017, 46, Basle: Bank for International Settlements.

Giordana, G. and M. Ziegelmeyer (2020): Stress testing household balance sheets in Luxembourg. Quarterly Review of Finance and Economics, 76, 115-138. Previous version published as BCL Working Paper 121. July 2018.

Girshina, A., T.Y. Mathä and M. Ziegelmeyer (2017): The Luxembourg Household Finance and Consumption Survey: Results from the 2nd wave, BCL Working Paper 106.

Grodecka, A. (2017): On the effectiveness of loan-to-value regulation in a multiconstraint framework, Sveriges Riksbank Working Paper No. 347.

Gross, M., and F.J. Población García (2017): Assessing the efficacy of borrower-based macroprudential policy using an integrated micro-macro model for European households. Economic modelling, 61, 510-528.

Hallessey, N., R. Kelly and T. O'Malley (2014): Macro-prudential tools and credit risk of property lending at Irish banks, Central Bank of Ireland Economic Letter Series, vol. 2014, No. 10.

Hanley, J.A., and B. McNeil (1982): The meaning and use of the Area under a Receiver Operating Characteristic (ROC) Curve. Radiology, 143(1), 29-36.

Household Finance and Consumption Network (2013): The Eurosystem Household Finance and Consumption Survey – Results from the First Wave. ECB Statistics Paper No. 2.

Household Finance and Consumption Network (2016): The Household Finance and Consumption Survey – results from the second wave. ECB Statistics Paper No. 18.

Household Finance and Consumption Network (2020): The Household Finance and Consumption Survey – results from the third wave. ECB Statistics Paper No. 36.

Hsieh, F., and B. W. Turnbull (1996): Nonparametric and semiparametric estimation of the receiver operating characteristic curve." Ann. Statist., 24(1), 25 - 40.

Jordà, Ò. (2012): Discussion of 'Anchoring countercyclical capital buffers: The role of credit aggregates'. International Journal of Central Banking, 7(4), 241-259.

Jordà, Ò., and A., Taylor (2011): Performance evaluation of zero net-investment strategies. NBER Working Paper, No 17150.

Kaminsky, G., Lizondo, S., and C. Reinhart (1998): Leading indicators of currency crises. IMF Staff Papers, 45(1), 1–48.

Kim, J., Choi, K., Kim, G. and Y. Suh (2012): Classification cost: An empirical comparison among traditional classifier, Cost-Sensitive Classifier, and MetaCost. Expert Systems with Applications, 39(4), 4013-4019.

Leika, M., and D. Marchettini (2017): A Generalized Framework for the Assessment of Household Financial Vulnerability. IMF Working Paper WP/17/228.

Makdissi, P., and Q. Wodon (2002): Consumption dominance curves: testing for the impact of indirect tax reforms on poverty. Economics letters, 75(2), 227–235.

Meriküll, J., and T. Rõõm (2020): Stress tests of the household sector using microdata from survey and administrative sources. International Journal of Central Banking, 16(2), 203-248.

Nier, E., Popa, R., Shamloo, M., and L. Voinea (2019): Debt Service and Default: Calibrating Macroprudential Policy Using Micro Data. IMF WP/19/182.

Sangaré, I. (2019): Housing sector and optimal macro-prudential policy in an estimated DSGE model for Luxembourg. BCL Working Paper 129.

9 Appendix

Appendix A: Tables and Figures for debt burden ratios

Table 12: Median debt burden ratios in 2018 for households with a recent HMR mortgage, by household characteristic

(In %)	Loan-to-value of HMR ratio		Mortgage debt-to-disposable income ratio		Total debt-to-disposable income ratio		Debt service-to-disposable income ratio		Mortgage maturity	
Variable	Median	Std. Err.	Median	Std. Err.	Median	Std. Err.	Median	Std. Err.	Median	Std. Err.
All households	88.6%	3.7%	643.6%	36.7%	647.0%	38.1%	34.5%	2.6%	25.0	0.5
Gender										
Female	98.9%	3.5%	757.9%	65.5%	764.6%	64.4%	39.1%	5.1%	25.0	2.2
Male	84.3%	6.3%	587.6%	44.3%	587.7%	46.3%	32.5%	3.0%	25.0	1.5
Age of reference person										
16-34	95.7%	5.0%	736.3%	69.8%	757.6%	72.6%	37.8%	4.9%	30.0	2.3
35-44	92.8%	9.7%	670.5%	67.9%	698.0%	65.4%	33.8%	3.9%	25.0	1.8
45-54	50.0%	12.2%	338.9%	28.3%	357.5%	47.8%	25.6%	6.0%	16.0	2.0
55-64	66.7%	20.2%	472.5%	165.4%	484.3%	168.6%	26.1%	11.6%	15.0	2.9
65+	80.0%	.	4199.1%	.	4199.1%	.	95.6%	.	10.0	.
Number of household members										
1	91.5%	7.8%	691.7%	84.9%	706.3%	84.2%	41.7%	5.6%	25.0	2.7
2	83.5%	12.2%	466.5%	74.8%	476.8%	59.0%	25.4%	5.0%	20.0	2.7
3	87.9%	7.9%	649.4%	102.9%	662.8%	102.6%	38.6%	4.5%	25.0	1.7
4	88.9%	9.3%	657.0%	82.9%	676.4%	89.7%	36.7%	5.0%	25.0	2.2
5+	87.9%	14.6%	728.7%	141.3%	742.9%	143.8%	28.8%	10.8%	30.0	2.8
Civil status										
Single	94.4%	6.5%	711.0%	84.0%	745.0%	89.6%	35.1%	4.5%	25.0	2.1
Couple	87.7%	5.6%	598.6%	49.4%	610.7%	53.4%	33.5%	3.5%	25.0	2.1
Divorced	91.4%	15.4%	664.0%	121.9%	664.7%	153.6%	36.3%	6.5%	22.0	3.2
Widowed	86.2%	3.2%	815.8%	1834.5%	815.8%	1834.5%	56.3%	32.9%	30.0	9.6
Country of birth										
Belgium	89.9%	27.1%	500.9%	171.1%	515.4%	168.7%	50.7%	19.8%	20.0	3.7
Germany	94.2%	16.9%	1510.7%	496.2%	1518.3%	498.3%	130.5%	41.9%	30.0	3.1
France	82.4%	19.7%	503.5%	74.0%	503.5%	73.1%	27.3%	4.6%	20.0	1.9
Italy	60.8%	19.6%	380.0%	182.6%	476.7%	184.8%	24.0%	9.6%	20.0	3.0
Luxembourg	86.4%	5.1%	640.2%	53.8%	643.9%	57.2%	33.1%	2.9%	25.0	1.1
Portugal	98.1%	14.1%	702.3%	89.1%	715.2%	87.5%	37.6%	9.0%	25.0	3.2
Other countries	99.3%	12.3%	879.2%	165.7%	879.2%	166.8%	44.8%	10.7%	30.0	2.8
Education level										
Low (ISCED=0:2)	49.3%	31.5%	397.5%	175.5%	466.7%	177.7%	44.2%	54.5%	15.0	4.6
Middle (ISCED=3,4)	92.7%	7.9%	698.9%	88.2%	717.8%	87.0%	33.6%	4.7%	25.0	1.9
High (ISCED=5,6 or 5:8)	88.2%	5.3%	619.8%	51.4%	628.5%	48.4%	34.5%	2.9%	25.0	1.4
Employment status										
Employed	87.5%	4.1%	650.7%	37.7%	659.9%	42.3%	35.2%	2.8%	25.0	0.9
Self-employed	92.2%	16.2%	509.1%	432.1%	520.0%	429.6%	33.6%	27.7%	25.0	4.9
Unemployed	116.0%	16.8%	728.7%	156.7%	739.6%	156.6%	28.2%	7.4%	30.0	2.4
Retired	66.7%	59.5%	243.3%	954.4%	243.3%	953.4%	20.6%	18.0%	20.0	5.4
Other	18.4%	79.0%	38.4%	1253.2%	38.4%	1253.2%	8.4%	38.2%	5.0	8.2
Housing status										
Owner-outright										
Owner with mortgage	88.6%	3.7%	643.6%	36.7%	647.0%	38.1%	34.5%	2.6%	25.0	0.5
Renter or other										
Gross Income quintile										
Q1	79.4%	21.0%	1122.7%	1014.0%	1127.4%	1015.4%	67.9%	59.5%	17.0	7.6
Q2	100.0%	20.3%	742.8%	187.9%	754.2%	190.8%	34.4%	9.6%	30.0	2.9
Q3	86.8%	10.5%	689.8%	75.6%	709.7%	70.0%	36.5%	6.6%	25.0	2.4
Q4	92.6%	10.4%	609.2%	63.9%	631.1%	66.0%	32.2%	5.2%	25.0	2.3
Q5	87.6%	5.3%	528.3%	71.4%	530.8%	69.4%	27.3%	3.4%	24.8	2.0
Net wealth quintile										
Q1	117.2%	14.9%	1043.4%	234.0%	1048.6%	233.7%	48.7%	13.5%	25.0	2.7
Q2	97.6%	3.3%	701.1%	55.7%	712.7%	51.3%	34.9%	3.6%	25.0	2.3
Q3	79.2%	11.7%	574.2%	135.0%	583.1%	127.7%	36.4%	5.4%	21.0	3.9
Q4	52.8%	15.6%	525.3%	222.2%	548.1%	191.1%	38.0%	29.6%	24.0	5.1
Q5	67.4%	11.7%	437.7%	111.6%	437.9%	111.3%	24.4%	6.1%	20.0	2.5

Source: Own calculations based on waves 1-3 of the LU-HFCS; data are multiply imputed and weighted; variance estimation based on 1000 replicate weights.

Table 13: Debt burden ratios at loan origination: households with recent HMR or OREP mortgages

Debt burden ratios	Wave 2010 Median	Wave 2010 Std. Err.	Wave 2014 Median	Wave 2014 Std. Err.	Wave 2018 Median	Wave 2018 Std. Err.	t-statistics of difference btw. 2018 - 2010	t-statistics of difference btw. 2018 - 2014
Initial loan-to-value (LTV) ratio of HMR	90.2%	7.1%	79.0%	4.5%	88.3%	3.8%	0.24	-1.57
Mortgage debt-to-disposable income (MDI) ratio	n.a.		558.9%	46.9%	602.5%	46.1%	n.a.	-0.64
Debt-to-disposable income (DI) ratio	n.a.		567.5%	47.9%	615.3%	47.2%	n.a.	-0.69
Debt service-to-disposable income (DSI) ratio	n.a.		34.5%	1.6%	33.1%	1.8%	n.a.	0.59
Mortgage maturity (MM)	25.00	1.62	25.00	2.17	25.00	1.23	0.00	0.00

Source: Own calculations based on waves 1st, 2nd and 3rd of the LU-HFCS; data are multiply imputed and weighted; variance estimation based on 1000 replicate weights. Results for the debt service-to-income ratio differ from those in Girshina, Mathä and Ziegelmeyer (2017) because leasing payments are now included. Income and debt are measured at time of latest mortgage origination. n.a. = net income not available in the 2010 wave.

Appendix B: Signals approach - robustness analysis

This appendix provides a robustness analysis of the signals approach exercise, using an alternative criterion to identify financially vulnerable households. The benchmark exercise in the main text focused on the impact of household financial vulnerability on banks, identifying vulnerable households as those with a loss given default (LGD) greater than zero. Instead, the robustness exercise focuses on those households that could potentially face problems servicing their debt, identifying vulnerable households as those with a probability of default (PD) greater than zero. PD>0 is a less restrictive criterion, since only some households with PD>0 have LGD>0, while all households with LGD>0 have PD>0 (see section 3).

In the 2018 baseline scenario, results when the condition variable is PD>0 are quite similar to the benchmark (LGD>0). The performance using individual ratios is comparable, except for the LTV ratio (Table 14). In addition, “optimal” limits are similar. The performance when combining ratios is similar to the benchmark but “optimal” limits are higher for the LTV ratio and lower for the other ratios (Table 16 and Table 17).

Figure 9 depicts the Receiver Operating Characteristic (ROC) curves for each debt burden ratio in the 2018 baseline. The ROC curve plots type I and type II classification errors for all the limits evaluated in the grid search. The y-axis reports the True Positive Rate (1 – type II error) and the x-axis reports the False Positive Rate (type I error). The origin represents the highest value considered for the limit (misses all vulnerable households but avoids misclassifying non-vulnerable households). Moving away from the origin, the value of the limit declines, reducing type II errors but raising type I errors. Therefore, the 45° line represents a poor performance (linear combinations of the maximum type II error at the origin with the maximum type I error at the top right corner).

In panel (a) of Figure 9 the condition variable is unity for households with PD>0. On this definition of financial vulnerability, the DI ratio (yellow), MDI ratio (purple) and DSI ratio (grey) deviate from the 45° line and therefore are effective at identifying financially vulnerable households. The LTV ratio (red line) performs more poorly, but the shape of the ROC curve provides some insights. For the LTV ratio, the ROC rises above the 45° line near the origin. Thus, as the limit is lowered from very high levels, type II error falls more rapidly than type I error increases. This suggests that high limits on the LTV ratio are effective at identifying financially vulnerable households. However, the ROC soon drops below the 45° line and does not recover in the upper quadrant, suggesting that at lower levels of the limit, type II error does not decline as rapidly as type I error increases. This suggests that, in the 2018 baseline many financially vulnerable households are characterized by relatively low LTV ratios. The MDI, DI and DSI ratios perform much better on type II error, meaning that they more rarely fail to identify vulnerable households.

Results are similar in panel (b) of Figure 9, in which the condition variable is unity for households with LGD>0. The LTV ratio performs better than with PD>0 in panel (a). The DI and MDI ratios appear to perform best, especially in terms of type II errors.

Figure 9: Receiver Operating Characteristic (ROC) curves of individual debt burden ratios in 2018 baseline

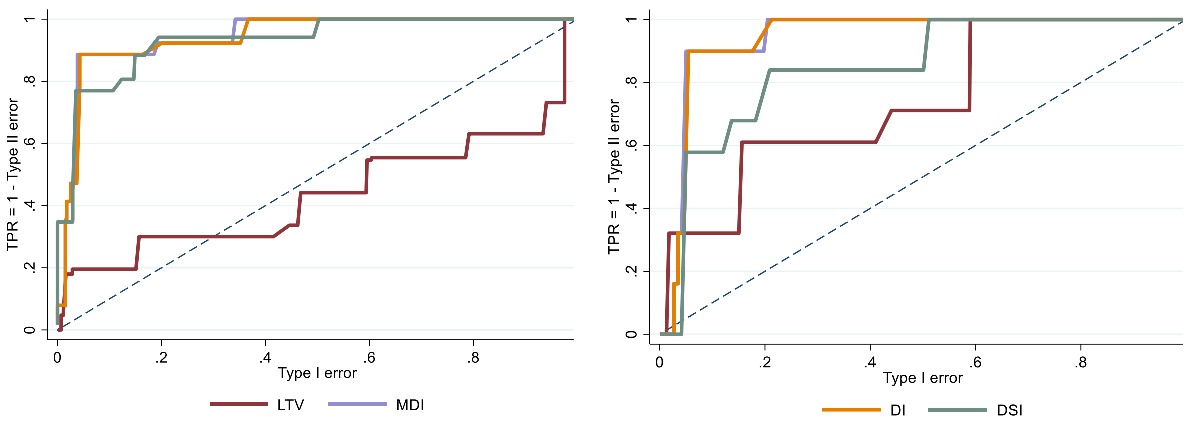

(a) Vulnerability definition I: PD>0 (b) Vulnerability definition II: LGD>0

Source: Own calculations based on the 3rd wave of the LU-HFCS; data are multiply imputed and weighted; the classification errors are calculated by pooling all the implicates.

In the adverse scenario, results change substantially. As in the benchmark exercise (LGD>0), using individual ratios is less effective than in the adverse scenario than in the 2018 baseline, with the exception of the LTV ratio, which outperforms all the other ratios in terms of AUROC (Table 15). The optimal limits for the LTV ratio are 175%, 95% and 75% depending on the loss function (θ equal to 0.25, 0.5 and 0.75, respectively). Policies that combine ratios are less effective than in the benchmark exercise (Table 18) and optimal limits are outside the legal range, in particular for the MDI ratio (Table 19).

Table 14: Performance of individual debt burden ratios in 2018 baseline (PD>0)

Fin. vuln. measure	Indic.	AUROC (s.d.)[a]	Weight on type II errors (θ)	Optimal Limit*	FPR	1-TPR	Loss	PPV	NPV	Marked.
PD>0	DI	0.921 (0.0038)	0.25	14.75	0.044	0.114	0.062	0.358	0.997	0.355
			0.5	14.75	0.044	0.114	0.079	0.358	0.997	0.355
			0.75	7.25	0.376	0.000	0.094	0.069	1.000	0.069
	MDI	0.920 (0.0038)	0.25	14	0.040	0.114	0.059	0.380	0.997	0.376
			0.5	14	0.040	0.114	0.077	0.380	0.997	0.376
			0.75	7.25	0.365	0.000	0.091	0.071	1.000	0.071
	DSI	0.871 (0.0047)	0.25	0.85	0.043	0.230	0.090	0.334	0.993	0.328
			0.5	0.54	0.195	0.058	0.127	0.119	0.998	0.117
			0.75	0.54	0.195	0.058	0.092	0.119	0.998	0.117
	LTV	0.425 (0.0056)	0.25	1.85	0.039	0.804	0.231	0.122	0.977	0.099
			0.5	1.85	0.039	0.804	0.422	0.122	0.977	0.099
			0.75	0.15	0.983	0.000	0.246	0.028	1.000	0.028

*Source: Own calculations based on the 3rd wave of the LU-HFCS; data are multiply imputed and weighted; the statistics are calculated by pooling all implicates. Notes: See Table 2 for definition of LTV, MDI, DI and DSI ratios and Table 6 for definition of FPR, TPR, Loss, PPV, NPV and Markedness. [a] Based on Hanley and McNeil (1982). * Maximum legal limits are 1 for the LTV ratio, 12 for the DI/MDI ratio, 0.75 for the DSI ratio and 35 years for MM (see Table 2).*

Table 15: Performance of individual debt burden ratios in adverse scenario (PD>0)

Fin. vuln. measure	Indic.	AUROC (s.d.)[a]	Weight on type II errors (θ)	Optimal Limit*	FPR	1-TPR	Loss	PPV	NPV	Marked.
PD>0	DI	0.607 (0.0037)	0.25	19.5	0.039	0.966	0.271	0.070	0.919	-0.011
			0.5	7	0.302	0.410	0.356	0.130	0.957	0.087
			0.75	2.25	0.008	0.931	0.239	0.085	0.990	0.076
	MDI	0.601 (0.0037)	0.25	9.25	0.181	0.525	0.267	0.187	0.947	0.134
			0.5	9.25	0.181	0.525	0.353	0.187	0.947	0.134
			0.75	2.25	0.920	0.015	0.242	0.086	0.983	0.069
	DSI	0.631 (0.0037)	0.25	0.5	0.204	0.385	0.249	0.209	0.959	0.168
			0.5	0.47	0.250	0.331	0.290	0.190	0.963	0.153
			0.75	0.15	0.921	0.000	0.230	0.087	1.000	0.087
	LTV	0.645 (0.0037)	0.25	1.85	0.045	0.970	0.276	0.055	0.918	-0.026
			0.5	0.95	0.387	0.295	0.341	0.137	0.960	0.097
			0.75	0.75	0.617	0.068	0.205	0.117	0.985	0.101

*Source: Own calculations based on the 3rd wave of the LU-HFCS; data are multiply imputed and weighted; the statistics are calculated by pooling all implicates. Notes: See Table 2 for definition of LTV, MDI, DI and DSI ratios and Table 6 for definition of FPR, TPR, Loss, PPV, NPV and Markedness. [a] Based on Hanley and McNeil (1982). * Maximum legal limits are 1 for the LTV ratio, 12 for the DI/MDI ratio, 0.75 for the DSI ratio and 35 years for MM (see Table 2).*

Table 16: Performance of combined debt burden ratios in 2018 baseline: DI, DSI, LTV ratios and MM

	Weight on type II errors (θ)	AUROC (s.d.)[a]	Rule	LTV	DI	DSI	MM	FPR	1-TPR	Loss	PPV	NPV	Marked.
PD>0	0.25	0.855 (0.005)	1	1.90	20.50	1.10	35	0.11	0.11	0.109	0.19	1.00	0.18
	0.5		1	1.90	17.46	0.57	35	0.20	0.00	0.102	0.12	1.00	0.12
	0.75		1	1.90	17.46	0.57	35	0.20	0.00	0.051	0.12	1.00	0.12
	0.25	0.903 (0.004)	2	1.80	14.42	0.86	35	0.03	0.11	0.049	0.48	1.00	0.47
	0.5		2	1.90	7.34	0.57	35	0.13	0.00	0.064	0.18	1.00	0.18
	0.75		2	1.90	7.34	0.57	35	0.13	0.00	0.032	0.18	1.00	0.18
	0.25	0.612 (0.006)	3	0.05	10.37	0.90	35	0.02	0.23	0.071	0.54	0.99	0.53
	0.5		3	0.15	14.42	0.25	35	0.07	0.11	0.092	0.26	1.00	0.26
	0.75		3	0.05	6.32	0.53	35	0.16	0.06	0.083	0.14	1.00	0.14
	0.25	0.382 (0.005)	4	0.05	11.64	0.95	5	0.02	0.50	0.138	0.43	0.99	0.42
	0.5		4	0.15	7.34	0.57	11	0.10	0.38	0.242	0.15	0.99	0.14
	0.75		4	0.15	7.34	0.57	11	0.10	0.38	0.313	0.15	0.99	0.14
LGD>0	0.25	0.857 (0.008)	1	1.90	20.50	1.10	35	0.12	0.10	0.116	0.07	1.00	0.07
	0.5		1	1.90	18.47	0.62	35	0.20	0.00	0.098	0.05	1.00	0.05
	0.75		1	1.90	18.47	0.62	35	0.20	0.00	0.049	0.05	1.00	0.05
	0.25	0.917 (0.006)	2	1.85	14.93	0.88	35	0.04	0.10	0.055	0.19	1.00	0.18
	0.5		2	0.92	13.41	0.62	35	0.12	0.00	0.058	0.08	1.00	0.08
	0.75		2	0.92	13.41	0.62	35	0.12	0.00	0.029	0.08	1.00	0.08
	0.25	0.800 (0.009)	3	0.82	14.93	0.27	35	0.05	0.10	0.061	0.16	1.00	0.16
	0.5		3	0.82	14.93	0.27	35	0.05	0.10	0.074	0.16	1.00	0.16
	0.75		3	1.80	9.36	0.25	23	0.16	0.00	0.040	0.06	1.00	0.06
	0.25	0.544 (0.010)	4	0.77	14.68	0.87	29	0.01	0.42	0.112	0.40	1.00	0.40
	0.5		4	0.82	8.35	0.62	23	0.07	0.32	0.193	0.09	1.00	0.09
	0.75		4	0.82	8.35	0.62	23	0.07	0.32	0.257	0.09	1.00	0.09

*Source: Own calculations based on the 3rd wave of the LU-HFCS; data are multiply imputed and weighted; the statistics are calculated by pooling all implicates. Notes: See Table 2 for definition of LTV, MDI, DI and DSI ratios and Table 6 for definition of FPR, TPR, Loss, PPV, NPV and Markedness. [a] Based on Hanley and McNeil (1982). * Maximum legal limits are 1 for the LTV ratio, 12 for the DI/MDI ratio, 0.75 for the DSI ratio and 35 years for MM (see Table 2).*

Table 17: Performance of combined debt burden ratios in 2018 baseline: MDI, DSI, LTV and MM

	Weight on type II errors (θ)	AUROC (s.d.)[a]	Rule	LTV	MDI	DSI	MM	FPR	1-TPR	Loss	PPV	NPV	Marked.
PD>0	0.25	0.856	1	1.90	20.50	1.10	35	0.11	0.11	0.108	0.19	1.00	0.19
	0.5	(0.005)	1	1.90	17.46	0.57	35	0.20	0.00	0.102	0.12	1.00	0.12
	0.75		1	1.90	17.46	0.57	35	0.20	0.00	0.051	0.12	1.00	0.12
	0.25	0.904	2	1.90	14.42	0.86	35	0.03	0.11	0.049	0.48	1.00	0.47
	0.5	(0.004)	2	1.90	7.34	0.57	35	0.11	0.00	0.057	0.20	1.00	0.20
	0.75		2	1.90	7.34	0.57	35	0.11	0.00	0.029	0.20	1.00	0.20
	0.25	0.608	3	0.10	13.92	0.43	35	0.03	0.17	0.069	0.40	1.00	0.39
	0.5	(0.006)	3	0.10	14.42	0.25	35	0.06	0.11	0.088	0.28	1.00	0.28
	0.75		3	0.05	6.32	0.53	35	0.16	0.06	0.083	0.14	1.00	0.14
	0.25	0.381	4	0.05	5.31	0.90	11	0.02	0.50	0.138	0.43	0.99	0.42
	0.5	(0.005)	4	0.15	7.34	0.57	11	0.10	0.38	0.241	0.15	0.99	0.14
	0.75		4	0.15	7.34	0.57	11	0.10	0.38	0.313	0.15	0.99	0.14
LGD>0	0.25	0.858	1	1.90	20.50	1.10	35	0.12	0.10	0.115	0.07	1.00	0.07
	0.5	(0.008)	1	1.90	18.47	0.62	35	0.20	0.00	0.098	0.05	1.00	0.05
	0.75		1	1.90	18.47	0.62	35	0.20	0.00	0.049	0.05	1.00	0.05
	0.25	0.918	2	1.90	14.93	1.08	35	0.04	0.10	0.055	0.19	1.00	0.18
	0.5	(0.006)	2	0.92	13.41	0.62	35	0.11	0.00	0.055	0.08	1.00	0.08
	0.75		2	0.92	13.41	0.62	35	0.11	0.00	0.028	0.08	1.00	0.08
	0.25	0.798	3	1.80	14.68	0.26	29	0.05	0.10	0.059	0.16	1.00	0.16
	0.5	(0.009)	3	1.80	14.68	0.26	29	0.05	0.10	0.073	0.16	1.00	0.16
	0.75		3	1.85	8.85	0.23	23	0.18	0.00	0.046	0.05	1.00	0.05
	0.25	0.544	4	0.82	4.04	0.85	29	0.01	0.42	0.112	0.40	1.00	0.40
	0.5	(0.001)	4	0.82	8.35	0.62	23	0.07	0.32	0.193	0.09	1.00	0.09
	0.75		4	0.82	8.35	0.62	23	0.07	0.32	0.257	0.09	1.00	0.09

*Source: Own calculations based on the 3rd wave of the LU-HFCS; data are multiply imputed and weighted; the statistics are calculated by pooling all implicates. Notes: See Table 2 for definition of LTV, MDI, DI and DSI ratios and Table 6 for definition of FPR, TPR, Loss, PPV, NPV and Markedness. [a] Based on Hanley and McNeil (1982). * Maximum legal limits are 1 for the LTV ratio, 12 for the DI/MDI ratio, 0.75 for the DSI ratio and 35 years for MM (see Table 2).*

Table 18: Performance of combined debt burden ratios in adverse scenario: DI, DSI, LTV and MM

	Weight on type II errors (θ)	AUROC (s.d.)[a]	Rule	LTV	DI	DSI	MM	1-TPR	FPR	Loss	PPV	NPV	Marked.
PD>0	0.25	0.614 (0.004)	1	1.90	16.45	0.53	35	0.42	0.22	0.270	0.19	0.95	0.15
	0.5		1	1.90	14.93	0.47	35	0.29	0.28	0.282	0.19	0.97	0.15
	0.75		1	0.98	19.99	0.88	23	0.02	0.73	0.199	0.11	0.99	0.10
	0.25	0.637 (0.004)	2	1.03	19.99	0.47	35	0.50	0.07	0.174	0.40	0.95	0.36
	0.5		2	1.85	19.99	0.47	23	0.35	0.14	0.243	0.30	0.96	0.27
	0.75		2	0.98	16.96	0.15	23	0.04	0.67	0.197	0.11	0.99	0.10
	0.25	0.664 (0.004)	3	1.03	19.99	0.47	23	0.50	0.03	0.150	0.57	0.96	0.52
	0.5		3	0.72	19.99	0.47	23	0.35	0.11	0.229	0.35	0.97	0.32
	0.75		3	0.98	1.26	0.13	23	0.04	0.67	0.197	0.11	0.99	0.10
	0.25	0.627 (0.004)	4	1.03	4.80	0.47	23	0.50	0.02	0.136	0.74	0.96	0.70
	0.5		4	0.72	4.80	0.47	23	0.35	0.09	0.219	0.40	0.97	0.37
	0.75		4	0.72	1.26	0.13	11	0.07	0.61	0.202	0.12	0.98	0.11
LGD>0	0.25	0.605 (0.005)	1	1.90	16.96	0.55	35	0.39	0.22	0.265	0.10	0.98	0.08
	0.5		1	1.90	20.50	0.49	35	0.23	0.28	0.253	0.10	0.99	0.09
	0.75		1	0.87	20.50	0.49	35	0.00	0.62	0.154	0.06	1.00	0.06
	0.25	0.663 (0.005)	2	1.03	18.47	0.62	35	0.57	0.04	0.176	0.29	0.98	0.26
	0.5		2	1.85	20.50	0.49	23	0.24	0.16	0.199	0.16	0.99	0.15
	0.75		2	0.87	20.50	0.49	23	0.00	0.48	0.119	0.08	1.00	0.08
	0.25	0.708 (0.005)	3	1.03	7.34	0.57	35	0.56	0.02	0.151	0.54	0.98	0.52
	0.5		3	0.56	20.50	0.49	23	0.24	0.13	0.181	0.20	0.99	0.19
	0.75		3	0.87	5.06	0.48	17	0.00	0.42	0.105	0.09	1.00	0.09
	0.25	0.671 (0.005)	4	0.72	5.31	0.49	23	0.24	0.11	0.138	0.23	0.99	0.22
	0.5		4	0.72	5.31	0.49	23	0.24	0.11	0.171	0.23	0.99	0.22
	0.75		4	0.72	4.55	0.26	17	0.00	0.43	0.107	0.09	1.00	0.09

*Source: Own calculations based on the 3rd wave of the LU-HFCS; data are multiply imputed and weighted; the statistics are calculated by pooling all implicates. Notes: See Table 2 for definition of LTV, MDI, DI and DSI ratios and Table 6 for definition of FPR, TPR, Loss, PPV, NPV and Markedness. [a] Based on Hanley and McNeil (1982). * Maximum legal limits are 1 for the LTV ratio, 12 for the DI/MDI ratio, 0.75 for the DSI ratio and 35 years for MM (see Table 2).*

Table 19: Performance of combined ratios in the adverse scenario: MDI, DSI, LTV and MM

	Weight on type II errors (θ)	AUROC (s.d.)	Rule	LTV	MDI	DSI	MM	1-TPR	FPR	Loss	PPV	NPV	Marked.
PD>0	0.25	0.617 (0.004)	1	1.90	16.45	0.53	35	0.42	0.22	0.270	0.19	0.95	0.15
	0.5		1	1.90	19.99	0.47	35	0.29	0.28	0.282	0.19	0.97	0.15
	0.75		1	0.98	16.96	0.96	23	0.02	0.73	0.199	0.11	0.99	0.10
	0.25	0.639 (0.004)	2	1.03	19.99	0.47	35	0.50	0.07	0.174	0.40	0.95	0.36
	0.5		2	1.90	19.99	0.47	23	0.35	0.14	0.243	0.30	0.96	0.27
	0.75		2	0.98	16.96	0.15	23	0.04	0.67	0.196	0.11	0.99	0.10
	0.25	0.660 (0.004)	3	1.03	19.99	0.47	23	0.50	0.03	0.150	0.57	0.96	0.52
	0.5		3	0.72	19.99	0.47	23	0.35	0.11	0.229	0.35	0.97	0.32
	0.75		3	0.98	1.26	0.13	23	0.04	0.67	0.197	0.11	0.99	0.10
	0.25	0.626 (0.004)	4	1.03	4.80	0.47	23	0.50	0.02	0.136	0.74	0.96	0.70
	0.5		4	0.72	4.80	0.47	23	0.35	0.09	0.219	0.40	0.97	0.37
	0.75		4	0.72	1.26	0.13	11	0.07	0.61	0.202	0.12	0.98	0.11
LGD>0	0.25	0.609 (0.005)	1	1.90	16.96	0.55	35	0.39	0.22	0.265	0.10	0.98	0.08
	0.5		1	1.90	20.50	0.49	35	0.23	0.28	0.253	0.10	0.99	0.09
	0.75		1	0.87	15.44	0.49	35	0.00	0.62	0.154	0.06	1.00	0.06
	0.25	0.664 (0.005)	2	1.03	18.47	0.62	35	0.57	0.04	0.176	0.29	0.98	0.26
	0.5		2	1.85	20.50	0.49	23	0.24	0.16	0.199	0.16	0.99	0.15
	0.75		2	0.87	20.50	0.49	23	0.00	0.48	0.119	0.08	1.00	0.08
	0.25	0.705 (0.005)	3	0.98	7.34	0.57	35	0.52	0.02	0.149	0.45	0.98	0.43
	0.5		3	0.62	20.50	0.49	23	0.24	0.13	0.181	0.20	0.99	0.19
	0.75		3	0.87	5.06	0.48	17	0.00	0.42	0.105	0.09	1.00	0.09
	0.25	0.670 (0.005)	4	0.72	5.31	0.49	23	0.24	0.11	0.138	0.23	0.99	0.22
	0.5		4	0.72	5.31	0.49	23	0.24	0.11	0.171	0.23	0.99	0.22
	0.75		4	0.72	4.55	0.26	17	0.00	0.43	0.107	0.09	1.00	0.09

*Source: Own calculations based on the 3rd wave of the LU-HFCS; data are multiply imputed and weighted; the statistics are calculated by pooling all implicates. Notes: See Table 2 for definition of LTV, MDI, DI and DSI ratios and Table 6 for definition of FPR, TPR, Loss, PPV, NPV and Markedness. [a] Based on Hanley and McNeil (1982). * Maximum legal limits are 1 for the LTV ratio, 12 for the DI/MDI ratio, 0.75 for the DSI ratio and 35 years for MM (see Table 2).*

Appendix C: Unemployment shock

The unemployment shock is likely to focus on households or individuals with certain characteristics. To allow for the non-uniform distribution of the shock, we estimate a logit model for the probability that an individual is unemployed using survey data for all household members who were part of the active work force[39] in the 2018 wave.

The vector of explanatory variables includes the following individual and household characteristics: gender, age, country of birth, marital status, highest educational attainment, household size, homeowner/tenant status, net wealth quintile. Explanatory variables also include the following characteristics that were only available in the 2018 wave: region dummies, current labour status or (if unemployed) previous labour status, language skills and current sector of employment (based on NACE codes) or previous sector of employment (if unemployed). Standard errors are clustered at the household level.

The unemployment rate observed across all individuals in the active work force covered by the 2018 HFCS wave was 4.7%.[40] After estimating the logit model, we adjust the intercept term to simulate the average unemployment rate increasing to 12%. Following Albacete and Fessler (2010) or Meriküll and Rõõm (2020), for each individual we draw a random number from the uniform distribution over the interval (0,1) and compare it to the logit-determined probability for the given individual to determine whether he or she becomes unemployed. We then adjust household net income as described in section 3.2 and recalculate the financial margin, probability of default and loss given default at the household level. Section 4.2.2 and 5 report estimated statistics for the adverse economic scenario that are averages across 1000 Monte Carlo iterations of this process. Each iteration requires simulating the employment status of every active individual.

[39] This includes all individuals between 16 and 64 years of age who reported their primary labour status as 1 - Doing regular work for pay / self-employed/working in family business; or 2 - On sick/maternity/other leave (except holidays), planning to return to work; or 3 - Unemployed.

[40] According to official figures (Statec Table B3019) the seasonally adjusted unemployment rate averaged 5.4% between April and November 2018 (when fieldwork took place for HFCS wave 3). This falls within our estimated confidence interval (3.8% to 5.6%).

Table 20: Logit coefficients of being unemployed

VARIABLES	Coef.	Std. Err.	
Gender: female [ref. category]*			
Gender*: male	-0.625	0.313	**
Age*	-0.101	0.086	
Age* squared	0.001	0.001	
Country of birth: Luxembourg [ref. category]*			
Country of birth*: Belgium	-0.368	0.982	
Country of birth*: Germany	2.474	0.680	***
Country of birth*: France	-0.983	0.803	
Country of birth*: Italy	-0.328	0.863	
Country of birth*: Portugal	-1.036	0.511	**
Country of birth*: other countries	-0.022	0.431	
Household size: 1 member [ref. category]			
Household size: 2 members	0.361	0.511	
Household size: 3 members	0.785	0.481	
Household size: 4 members	0.117	0.623	
Household size: 5+ members	1.108	0.520	**
Marital status: single/ widowed [ref. category]*			
Marital status*: couple	-0.464	0.343	
Marital status*: divorced	0.086	0.493	
Education level: Low (ISCED=0,1,2) [ref. category]*			
Education level*: Middle (ISCED=3,4)	-0.295	0.362	
Education level*: High (ISCED=5,6)	-0.445	0.449	
Housing status: owner-outright [ref. category]			
Housing status: owner-with mortgage	-0.412	0.398	
Housing status: renter or other	-0.170	0.605	
Net wealth quintile 1 [ref. category]			
Net wealth quintile 2	-0.318	0.506	
Net wealth quintile 3	-0.464	0.572	
Net wealth quintile 4	-1.647	0.805	**
Net wealth quintile 5	-2.584	0.826	***
Region: Capellen & Mersch [ref. category]			
Region: Esch-sur-Alzette	1.182	0.680	
Region: East	1.321	0.746	*
Region: Luxembourg	1.316	0.692	*
Region: North	0.933	0.684	
Current or past labour status: self-employed & other [ref. category]*			
Current or past labour status*: employee	-2.111	0.481	***
Language skills*: Luxembourgish perfect	-0.636	0.415	
Language skills*: French perfect	-0.107	0.440	
Language skills*: German perfect	-1.492	0.546	***
Language skills*: English perfect	1.272	0.571	**
Sector of employment: Non-market services (excl. public sector) (NACE: L-U excluding G) [ref. category]*			
Sector of employment*: Public sector (NACE: O P)	-1.650	1.132	
Sector of employment*: Financial services (NACE: K)	-0.438	1.042	
Sector of employment*: Wholesale and retail trade (NACE: G)	1.728	0.525	***
Sector of employment*: Agriculture, industry, and construction (NACE: A-F)	2.040	0.496	***
Sector of employment*: Market services (NACE: H I J)	0.992	0.489	**
Constant	0.351	1.968	
Observations	2,085		
Pseudo R2 (mean across implicates)	0.237		

Source: Own calculations based on the LU-HFCS wave 3. Results are imputed and weighted.
Note: Personal characteristics indicated by an asterisk.

BANQUE CENTRALE DU LUXEMBOURG
EUROSYSTÈME

2, boulevard Royal
L-2983 Luxembourg

Tél.: +352 4774-1
Fax: +352 4774 4910

www.bcl.lu • info@bcl.lu