DNB Working Paper

No. 635 / May 2019

Heterogeneity and Asymmetric Macroeconomic Effects of Changes in Loan-to-Value Limits

Jasper de Jong and Emmanuel de Veirman

DeNederlandscheBank

EUROSYSTEEM

Heterogeneity and Asymmetric Macroeconomic effects of changes in Loan-to-Value limits

Jasper de Jong and Emmanuel de Veirman *

* Views expressed are those of the authors and do not necessarily reflect official positions of De Nederlandsche Bank.

Working Paper No. 635

De Nederlandsche Bank NV
P.O. Box 98
1000 AB AMSTERDAM
The Netherlands

May 2019

Heterogeneity and Asymmetric Macroeconomic Effects of Changes in Loan-to-Value Limits*

Jasper de Jong and Emmanuel De Veirman†

May 2019

Abstract

We estimate the macroeconomic effects of changes in loan-to-value limits using an approach that involves the cross-sectional loan-to-value distribution and does not require that a limit is actually in place. We show that the effects are asymmetric and non-linear as tighter limits constrain a larger fraction of borrowers. Symmetry is a good approximation when the limit is tight but not at other points. We find that an increase in heterogeneity can substantially increase the effects of a change in loan-to-value caps. We document that if one abstracts from borrower heterogeneity, one understates the size of the effects of LTV limits when the limit lies above the average LTV.

JEL codes: C32, E21, E32, E44

Keywords: Credit constraints, macroprudential policy, loan-to-value ratio

*We thank Maurice Bun, Dorinth van Dijk, Machiel van Dijk, Ferre De Graeve, Leo de Haan, Jakob de Haan, Frank van Hoenselaar, Jesper Lindé, Remco van der Molen, Mauro Mastrogiacomo, Kostas Mavromatis, Mark Mink, Rob Nijskens, Frank Smets, Raf Wouters and Francesco Zanetti for input. This paper uses data from the CentERpanel and from the LISS (Longitudinal Internet Studies for the Social Sciences) panel administered by CentERdata (Tilburg University, the Netherlands). The views expressed in this paper are those of the authors and do not necessarily reflect those of De Nederlandsche Bank. Any errors are our own.

†De Veirman is a researcher at De Nederlandsche Bank. De Jong was involved in this project while at De Nederlandsche Bank. Corresponding author: `Manu.Veirman@dnb.nl`, P.O. Box 98, 1000 AB Amsterdam, the Netherlands.

1 Introduction

In this paper, we develop and implement a new approach to estimating the macroeconomic effects of changes in the regulatory maximum for loan-to-value ratios (LTVs) on mortgage loans.

Internationally, LTV limits are increasingly prevalent.[1] As a macroprudential policy instrument, LTV caps are typically imposed or changed so as to enhance the resilience of the financial system to shocks or to reduce the likelihood of such adverse shocks by containing house price and mortgage credit booms. However, it is not straightforward to estimate the effects of such limits, both because typically there is no sufficiently long time series for the LTV cap and because the cap typically varies infrequently over time.[2]

Faced with this challenge, regression analysis in the literature regarding the macroeconomic effects of LTV policy often consists of panel regressions pooling a range of economies and typically uses coded variables to indicate, for instance, whether an LTV cap is present or not.[3] Standard panel regressions abstract from differences across economies in terms of the effects of LTV actions. The use of coded variables imposes the restriction that different LTV policy actions have the same effect, independent of the size of the change in the limit.

Our approach to estimating the effects of LTV limits is as follows. We first use the cross-sectional LTV distribution to track how changes in the LTV limit affect the average of the LTV distribution. We then feed these changes in the average LTV as impulses into a Vector Error

[1] See, for instance, the databases on the use of macroprudential policy instruments including LTV limits that accompany Shim, Bogdanova, Shek, and Subelyte (2013) and Cerutti, Claessens, and Laeven (2017), as well as the appendix of Akinci and Olmstead-Rumsey (2017).

[2] Crowe et al. (2013) similarly note that among countries for which there is more than one data point for the LTV limit available in their database, there are very few in which the limit actually varies over time.

[3] For instance, studies on the effects of macroprudential policy such as Kuttner and Shim (2016), Cerutti, Claessens, and Laeven (2017) and Akinci and Olmstead-Rumsey (2018) all measure LTV policy using coded variables. See Galati and Moessner (2017) for a detailed review of the literature on the effects of macroprudential policy, including LTV policy. See Igan and Kang (2011) and Wong et al. (2011) for studies that measure the effects of LTV policy in particular.

Correction Model (VECM) and track the dynamic responses of house prices and aggregate output.

We illustrate our approach by means of the LTV policy action in the Netherlands that consisted of imposing an LTV cap in mid-2011 and lowering it, from 2013 through 2018, in annual steps of one percentage point. This gradual approach, which stands out internationally, is interesting: by comparing LTV distributions in nearby spells against the background that the cap differs by one percentage point, we build intuition on how changes in LTV caps affect the distribution.

Our approach does account for the actual size of changes in the LTV cap. We do not impose any cross-country coefficient restrictions. Our approach does not require that an LTV cap is actually in place, and can therefore be used for ex ante as well as ex post analysis.

To our knowledge, our paper is the first that relates heterogeneity to non-linearity and asymmetry in the effects of changes in LTV caps.

Our results are as follows.

First, our estimates suggest that loan-to-value policy actions in the Netherlands in 2011-2018 are having a non-trivial effect on relative house prices and aggregate real output. When using a comparable policy scenario, the estimated size of the effects is similar to that in our contribution to De Nederlandsche Bank (2015a). Caveats include the fact that if households carried out high-LTV purchases in anticipation of the imposition of the limit, that would tend to imply that our procedure for translating changes in the limit into changes in the average LTV overstates the effect. In addition, the standard caveat of Vector Autoregression and VECM analysis applies that the impulse responses do not prove economic causality.

Our study does not yield conclusions about which LTV policy would be optimal. The reasons for this include the fact that we do not analyze welfare effects. Furthermore, we focus on macro effects and do not consider other effects of LTV limits such as the effects on financial stability. See

CPB Netherlands Bureau for Economic Policy Analysis (2015, 2017) and De Nederlandsche Bank (2015a) for discussion and analysis of a more comprehensive set of plausible effects of a particular LTV policy scenario in the Netherlands.

Second, we find that changes in the LTV cap have non-linear effects in the sense that the effects become larger as the cap tightens. This occurs because the fraction of households for which the cap binds increases as the cap tightens.

Third, we document that a moderate, mean-preserving change in the degree of heterogeneity can substantially alter the effects of changes in LTV caps. We show that if the standard deviation of the LTV distribution had been 25 percent larger than in the data, the macroeconomic effects of the forementioned Dutch policy action would have been over twice as large as if the standard deviation had been 25 percent smaller than in the data.

Fourth, from a given initial level for the LTV cap and for a given underlying LTV distribution, a tightening tends to affect more borrowers than an easing, and therefore tends to have stronger effects. This relates to earlier findings that a typical tightening in the LTV cap has stronger macroeconomic effects than a typical easing.[4] Our approach differs from earlier papers in that asymmetry as well as the size of the effects is conditional on the relative position of the LTV cap and the LTV distribution. Tightenings have smaller effects when they occur from a point where the cap is binding for few households, while loosenings have larger effects when they occur from a point where many households are constrained by the cap.

Fifth, we find that with the LTV distribution that prevailed in the Netherlands, symmetry is a fairly good approximation for LTV limits around 80 percent or lower since in that area the limit

[4] Igan and Kang (2011) and Kuttner and Shim (2016) find that typically, tightenings in the LTV cap have larger effects than loosenings. McDonald (2015) finds that this also applies conditional on a credit expansion. Elliott, Feldberg, and Lehnert (2013) and Zdzienicka et al. (2015) document asymmetry in the same direction for US macroprudential policy. Each of these papers uses coded variables and measures a single effect for tightenings and a separate one for easings.

is binding for the bulk of households. Asymmetry prevails for higher LTV limits.

Sixth, when we impose homogeneity, there is extreme asymmetry for a narrow range of initial LTV limits, whereas asymmetry is more spread out in a heterogeneous world. This implies that when the cap exceeds the average LTV in a heterogeneous setting, the effects of changing the limit are larger than if all borrowers have the same LTV.

Our findings are relevant for the literature on credit constraints, in the sense that an economy-wide LTV limit is one particular type of credit constraint. In one class of New Keynesian general equilibrium models with a credit constraint that applies for borrowing households, known as "Two Agent New Keynesian" (TANK) models, the LTV is homogeneous in the sense that all borrowers are at the same LTV. This implies either that all responses are symmetric or, in most homogeneous-borrower models with occasionally binding credit constraints, that there is threshold asymmetry around the point where the credit constraint ceases to be binding.[5] Another class of models, referred to as Heterogeneous Agent New Keynesian (HANK) models, feature a net worth distribution induced by uninsurable idiosyncratic income shocks, as well as a constraint on the extent to which net worth can turn negative, which is interpreted as a borrowing limit. We are not aware of any paper investigating asymmetry and non-linearity in response to changes in the borrowing limit within a heterogeneous agent model.[6]

Our evidence also relates to a literature which examines the empirical relation between LTV ratios and house prices in the United States.[7] Our study differs from those studies in that we

[5]Symmetry applies in linearized models with constraints that are binding at all points of analysis, including the seminal paper by Iacoviello (2005). Threshold asymmetry in response to housing and LTV shocks applies in most models with occasionally binding constraints, including Justiniano, Primiceri, and Tambalotti (2015) and Guerrieri and Iacoviello (2017). Brzoza-Brzezina, Gelain, and Kolasa (2014) built a model with occasionally binding credit constraints where there is smooth asymmetry in response to LTV shocks.

[6]See Kaplan and Violante (2018), and the references therein, for HANK models. Specifically, Guerrieri and Lorenzoni (2017) examine the effects of a one-time tightening in the borrowing limit in a HANK framework. Debortoli and Gali (2017) compare HANK and TANK models in terms of their implications for the macroeconomic effects of macroeconomic shocks.

[7]See Duca, Muellbauer and Murphy (2011, 2012) and Glaeser, Gottlieb and Gyourko (2011).

document that heterogeneity gives rise to non-linearity and asymmetry, and in that we examine the effects of LTV limits. In the United States, changes in Federal Housing Administration downpayment requirements as well as LTV limits used by Government Sponsered Enterprises affect the distribution of LTVs, with the added complexity that different systems apply to conforming and non-conforming mortgages, a complexity which we do not have to deal with in our setting.

To construct our LTV series, we designed a survey and fielded it in two CentERdata household panels. Our survey is one-off but pertains to respondents' first home, so we obtain a series for the average LTV for first-time buyers in the Netherlands.[8] As we use a one-off survey, our series is not subject to any changes in methodology, unlike what may be the case if one combines evidence from various waves of a survey as would be the case with the American Housing Survey (AHS). This is relevant in that the leading study constructing and using a survey-based LTV series for the United States[9] uses the AHS. Another advantage is that we measure the LTV for respondents' first home irrespective of whether they still live in it or not. In contrast, the AHS has a flag for first-time buyers which selects respondents who still live in their first home, which plausibly results in selection bias in the sense that houses that are less frequently sold are more likely to still be in the sample.

Credit constraints are typically not among the fundamentals that determine house prices in standard real user cost models. In that sense, our result that Dutch loan-to-value limits had a non-trivial effect on house prices is in line with earlier papers that highlight the role of nonfundamental drivers of house prices.[10]

We do not examine the effects of LTV policy on financial stability or on welfare. We do

[8] This approach was inspired by Timmermans (2012).
[9] Duca, Muellbauer and Murphy (2011).
[10] See Ling, Ooi, and Le (2015) and the references therein.

not assess the effects of LTV limits on the strength of debt-collateral circles.[11] Unlike in the optimizing macro models which we referenced, we do not account for expectational effects such as the expectation that a collateral constraint is binding in the future.

Section 2 first shows how the Dutch LTV distribution changed when the cap changed and then explains our approach for translating changes in the LTV limit into changes in the average LTV. Section 3 discusses the time series data with particular emphasis on our survey-based measure for the average LTV. Section 4 estimates the macroeconomic effects of changes in the Dutch LTV cap. Section 5 examines the role of the degree of heterogeneity for the size of these effects. Section 6 investigates non-linearity and asymmetry. Section 7 concludes.

2 Capturing the Effect of the Cap on the Loan-to-Value Distribution

In this section, we first document how the empirical loan-to-value distribution evolved in the Netherlands as the loan-to-value limit was gradually tightened. Next, we explain how we use these loan-to-value distributions to translate changes in loan-to-value limits into changes in the typical loan-to-value ratio.

2.1 Successive Loan-to-Value Distributions and the Tightening Cap

Since August 1, 2011, Dutch mortgage providers have been subject to an LTV limit. Starting from that date through December 31, 2012, a new mortgage loan could effectively amount to at most 106 percent of the house price excluding sales taxes. The LTV limit was subsequently reduced by one percentage point each year through 2018. For instance, the LTV limit was 105 percent throughout the calendar year 2013 and it is 100 percent from January 1, 2018 onwards.[12]

[11] For studies that do, see Lamont and Stein (1999), Gerlach and Peng (2005) and Almeida, Campello, and Liu (2006).

[12] The LTV limit for the period August-2011-December 2012 is contained in a code of conduct for mortgage providers referred to as the “Gedragscode Hypothecaire Financiering 2011”. Compliance was subject to supervision

The fact that the changes in the LTV cap have been small and frequent stands out in an international perspective. These features are interesting for our purposes in that they allow us to obtain a relatively clean view on how changes in LTV limits affect actual LTV ratios. In particular, we compare empirical cross-sectional LTV distributions in successive spells, against the background that the LTV cap in two successive spells differs by one percentage point. That evidence informs our approach for measuring the ceteris paribus effect of changes in the cap, which we detail in Section 2.2.

To document successive LTV distributions at different levels of the cap, we use micro data on mortgage loans from the Loan Level Data (LLD) database. De Nederlandsche Bank (DNB) collects these confidential data from Dutch mortgage providers on behalf of Statistics Netherlands. Mortgage providers are asked to report their entire mortgage portfolio. The description of the LLD in De Nederlandsche Bank (2015b) indicates that in 2012Q4-2014Q4, the LLD covered between 77 and 81 percent of total outstanding Dutch mortgage debt.[13]

As the name of the database indicates, data are at the level of an individual loan. We use an LTV measure that aggregates individual loans when there are multiple loans associated with the same house purchase. The LTV is expressed as a ratio of the house price excluding closing costs paid for by the buyer such as sales taxes and notary fees. The LTV measure we use is meant to capture the LTV at the time of purchase of the property.[14]

by the "Autoriteit Financiële Markten", the entity responsible for behavioral supervision in the Netherlands. The LTV limits from January 1, 2013 onwards are provided for by a law enacted in December 2012. We treat the code of conduct and the law as a single policy action.

[13]In order to construct the LLD, De Nederlandsche Bank requests data regarding mortgages from the same financial intermediaries that report information to the European Central Bank (ECB) in order to apply for the ECB's acceptance of their securitized mortgages as collateral. As of 2012Q4, 9 Dutch banks and 3 Dutch insurance companies reported data for the LLD. Small banks, pension funds, foreign intermediaries and insurance companies that do not securitize mortgages do not report data for the LLD. This and more detailed information on the LLD is in De Nederlandsche Bank (2015b).

[14]We use the LLD variable called the "original LTV". This is the LTV at the time of mortgage origination to the extent that the recorded origination date actually marks an origination and not a loan renegotiation. We take the following steps to clean our data from observations which are likely to correspond to renegotiations. First, as we explain later in this subsection, we focus on young borrowers. Second, we drop observations for which the

We drop observations for which the LTV measure exceeds 200 percent. Because a house purchase only enters the LLD if it is associated with a mortgage loan, observations with zero LTVs should be spurious. We therefore drop any zero LTVs. We drop LTVs for which the total loan amount exceeds one million euro.[15]

The LLD consists of a succession of quarterly vintages of data which properly start in 2013Q1 (with a test wave in 2012Q4). Each vintage contains information for mortgage loans that are in a mortgage provider's portfolio in that quarter. In principle, we use the LLD vintage for the first quarter of year $t + 1$ to collect loans that originated in any calendar year t. In choosing the first available vintage after the completion of year t, our aim is to wait until all year-t loans are reported while keeping the number of loans that already dropped from the books to a minimum. For instance, we obtain loans originated in 2012 from the 2013Q1 LLD vintage. For mortgages dated 2010 and 2011, we use the first available Q1 vintage, i.e. 2013Q1.

We wish to select first-time buyers. One reason for doing so is that for repeat buyers, the LTV at the time of a repeat home purchase may be affected by the extent and sign of previous housing equity accumulation. However, the LLD does not indicate whether the borrower is a first-time buyer. It does, however, indicate the borrower's birth year. We therefore treat young buyers as a proxy for first-time buyers. We select borrowers who are 35 years or younger at the end of the year of mortgage origination.

Figures 1A through 1C show cross-sectional LTV distributions for this definition of young

reported appraisal date lies more than two years before the reported origination date. Third, we drop observations for which the reported origination date is less than 29 years before the reported maturity date. We do this against the background that in the years for which we use LLD data (2010-2017), on the order of 90 percent of new Dutch mortgages had a maturity duration of 30 years, with much of the remainder being interest-only loans for which the recorded maturity duration typically exceeds 30. As a separate issue, we drop observations for which the reported appraisal date lies after the origination date.

[15] For loans issued in 2010-2012, there is a single mortgage provider which reports a substantial mass of observations with LTVs at levels that we deem to be implausibly high. This suggests that it used different measurement practices in those years, for instance in terms of house price valuation. We drop loans originated by that provider from the sample irrespective of the origination date.

borrowers for different periods depending on the mortgage origination date. Every panel compares the distribution before and after a change in the cap, with blue bars indicating the earlier period and pink bars indicating the later period. The LTV is on the horizontal axis and the bin width is one percentage point. In all histograms in this paper, integer values fall at the right edge of any bin. For example, an LTV of 106 percent falls in the 105-106 bin. Bars indicate the share of observations in each bin.

The top panel of Figure 1A compares the distribution in the last twelve months before the introduction of the LTV limit (2010M8-2011M7, which we refer to as the “pre-cap” period) to that for the spell for which the LTV cap of 106 percent was in place (2011M8-2012M12, the “106 cap” period). Before the correction for tax changes which we carry out in Section 2.2, the mean of the 56,435 observations that constitute the pre-cap distribution is 94.46 percent. Note that a substantial fraction of observations exceeds 100 percent, while there are few observations exceeding 112 percent. This is consistent with the fact that it is fairly common in the Netherlands for buyers to use mortgage borrowing towards funding buyer-incurred closing costs. For most of the pre-cap period, sales taxes were at six percent, and total buyer-incurred closing costs typically amounted to 10 or 11 percent of the house price. In addition to sales taxes, these consist of notary and broker fees as well as payment of the premium for a government-backed mortgage default guarantee scheme known as the Nationale Hypotheek Garantie (NHG).

Figure 1A: Successive Loan-to-Value Distributions in the Netherlands

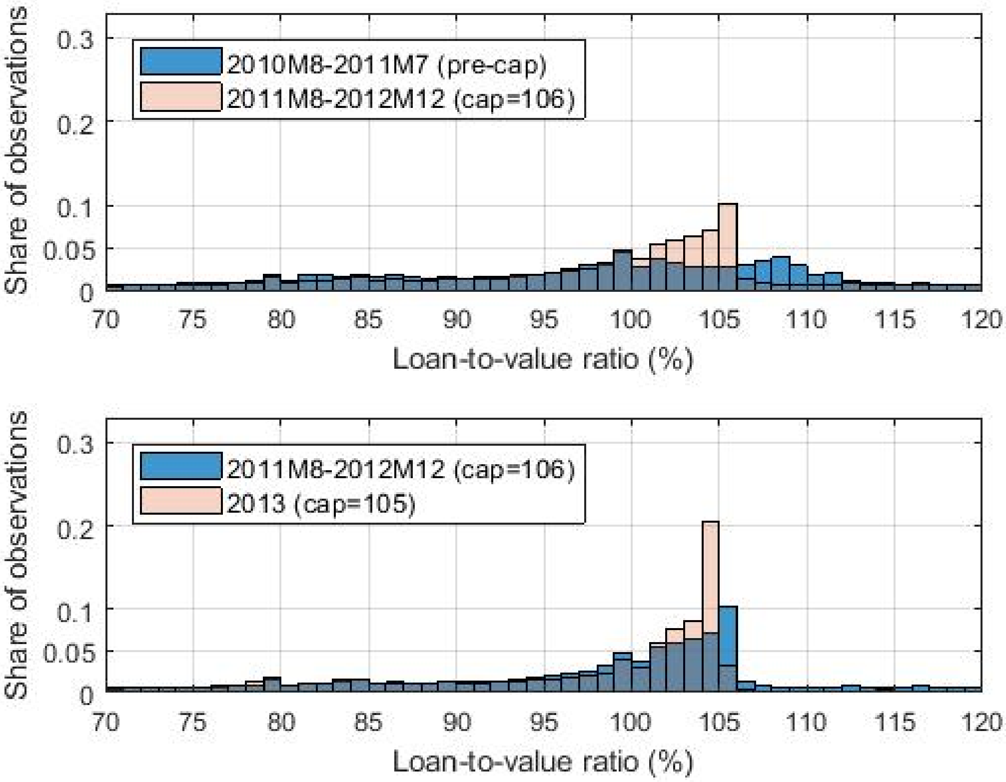

Note: This figure plots distributions for the Netherlands of the mortgage loan-to-value ratio for borrowers that were thirty-five years of age or younger at the end of the year of mortgage origination. Each panel compares the distribution in two successive regimes for the LTV limit. Each time, the blue distribution is the earlier period and the pink distribution is the later period. The top panel compares the distribution for 2010M8-2011M7, which were the last twelve months before the introduction of the LTV limit, to the distribution for 2011M8-2012M12, when the limit was 106 percent. The data in this figure are not adjusted for the change in sales taxes in June 2011. The bottom panel compares the distribution for the spell with a cap of 106 percent to that for the calendar year 2013, when the cap was 105 percent. Data are from the Loan Level Data (LLD) database, collected by De Nederlandsche Bank on behalf of Statistics Netherlands.

Figure 1B: Successive Loan-to-Value Distributions in the Netherlands

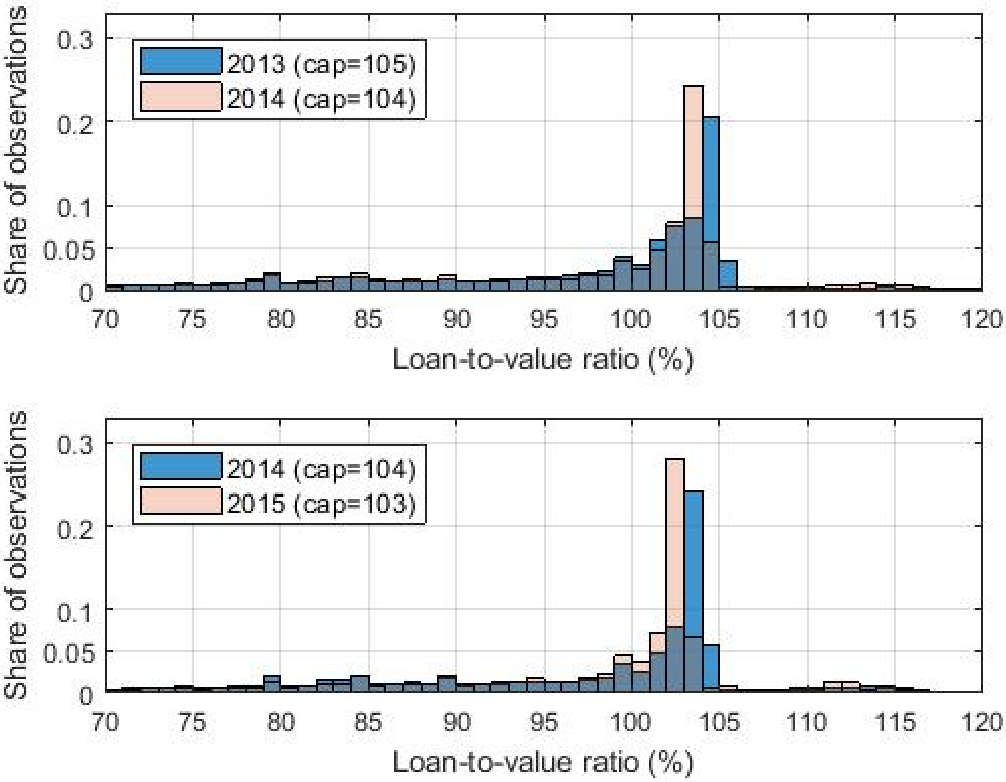

Notes are as under Figure 1A, except for the fact that the top panel compares the LTV distributions for 2013 and 2014 against the background of a reduction in the LTV limit from 105 percent to 104 percent, while the bottom panel compares the distributions for 2014 and 2015 against the background of a reduction in the limit from 104 to 103 percent.

Figure 1C: Successive Loan-to-Value Distributions in the Netherlands

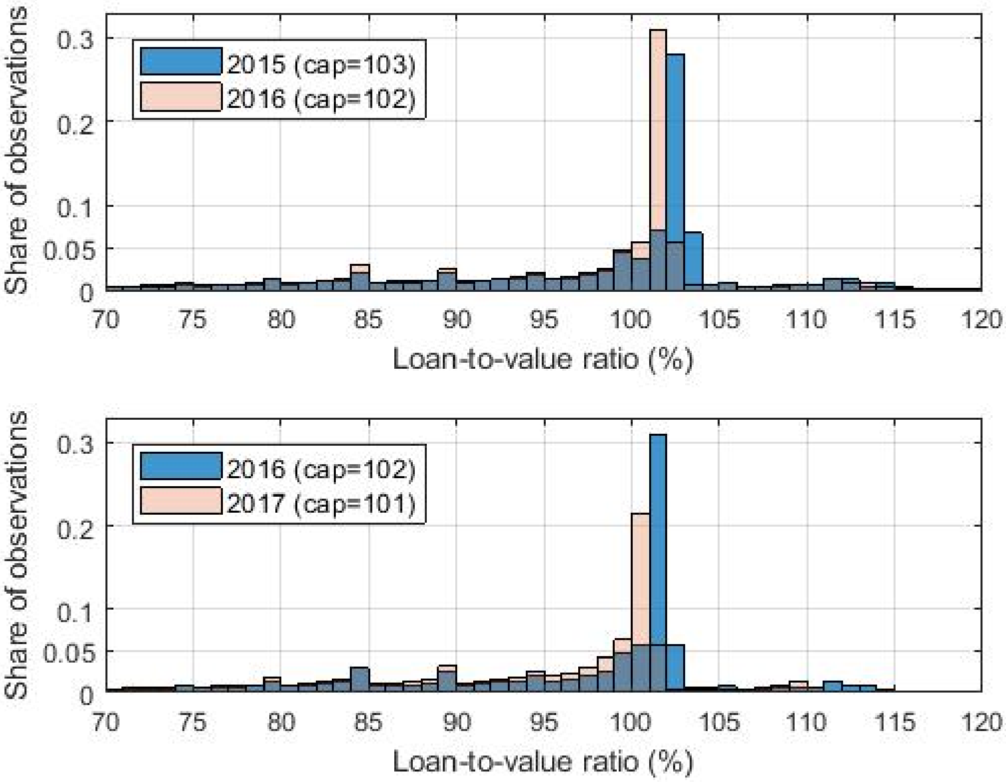

Notes are as under Figure 1A, except for the fact that the top panel compares the LTV distributions for 2015 and 2016 against the background of a reduction in the LTV limit from 103 percent to 102 percent, while the bottom panel compares the distributions for 2016 and 2017 against the background of a reduction in the limit from 102 to 101 percent.

Relative to the pre-cap distribution, the distribution in the 106-cap period reveals a decrease in the share of observations in the 106-112 percent range and an increase in the share of observations in the 100-106 percent range. The fact that the 106-cap distribution is virtually truncated at 106 suggests that this shift is in part due to the imposition of the cap at 106. Another feature which plausibly contributed to the declines in LTVs in that range is a cut in the tax on housing sales which we address in Section 2.2.

Each of the remaining panels of Figures 1A through 1C compares two successive LTV distributions against the background of a reduction in the LTV limit by one percentage point, through the reduction to 101 percent in 2017. For every distribution in which a cap is present, the distribution is virtually truncated at the cap.

Figures 1A through 1C do not reveal pronounced recurring patterns in the year-on-year changes in the mass of LTVs below the cap. Therefore, we see little evidence to suggest that changes in the cap alter the LTV of borrowers below the cap. Such an effect on unconstrained borrowers could in principle occur if a change in the cap is associated with a change in interest rates which affects borrowing through intertemporal substitution. As another example, borrowers may have a precautionary motive to borrow less than the limit, implying that a change in the cap could alter the LTV of such borrowers even if they are not actually constrained by the limit.[16]

Taking all six panels together, the distributions suggest that the effect of the cap on the distribution can be fairly well captured by setting the LTV for observations above the cap equal to the value of the cap. This is the essence of our approach in Section 2.2.

Still, at all levels of the cap, a non-trivial share of observations remains above the cap. This plausibly reflects the fact that the LTV limit does not apply to all borrowers. First-time borrowers can exceed the LTV limit if they use the mortgage loan to finance energy-saving home improvements or if their debt-service-to-income ratio is sufficiently low. As we explain in Section 2.2, our baseline approach accounts for exceptions.[17]

[16] Several panels in Figures 1A through 1C do reveal an increase in the mass of borrowers in a range of up to a few percentage points below the cap. This is quite pronounced in the upper panel of Figure 1A, but that reflects a decrease in the sales tax as well as the imposition of the cap. For the remaining five panels, any increase in the mass of borrowers close to but below the limit is small. We choose to treat it as not being an effect of the limit. Within a Heterogeneous Agents New Keynesian model, Guerrieri and Lorenzoni (2017) analyze the intertemporal substitution and precautionary saving effects across the wealth distribution of a tightening in the borrowing constraint.

[17] We do not take a stance on the actual empirical share of exceptions in the sense that some of the reported LTVs above the cap may not actually reflect exceptions for first-time buyers. In particular, we cannot rule out the possibility that some transactions in our sample are actually renegotiations on loans that originated during an earlier (and less strict) LTV regime. Furthermore, it is likely that our sample of young borrowers includes at least some repeat buyers, for whom additional exceptions apply, including the option for borrowers with negative equity to exceed the LTV limit. However, when we select very young borrowers (using 25 instead of 35 as the age criterion), there is still a substantial fraction of observations above the cap, which suggests that renegotiations and repeat buyers do not constitute the main reasons for LTVs above the cap. For a short overview of the exceptions to the LTV limit, see the following Dutch-language webpage:

2.2 Capturing the Effect of the Cap on the Distribution

Before measuring the effect of LTV limits, we correct the pre-cap distribution for the fact that the tax on housing sales declined from 6 to 2 percent effective on June 15, 2011. By doing so, we mean to avoid overstating the effect of the nearly coincident imposition of the LTV cap on the LTV distribution. For LTV observations with origination month May 2011 or earlier,[18] we compute LTV_{it}, which is the tax-corrected LTV, from $\widetilde{LTV}_{it}$, which is the unadjusted LTV, as follows:

$$\begin{cases} LTV_{it} = \widetilde{LTV}_{it} & \text{if } \widetilde{LTV}_{it} \leq 100 \\ LTV_{it} = \widetilde{LTV}_{it} - 4\left(\frac{\widetilde{LTV}_{it} - 100}{10}\right) & \text{if } 100 < \widetilde{LTV}_{it} < 110 \\ LTV_{it} = \widetilde{LTV}_{it} - 4 & \text{if } \widetilde{LTV}_{it} \geq 110 \end{cases} \tag{1}$$

This scheme is a proxy for the effect of sales taxes on the LTV distribution. It reflects our intuition that households may borrow less when the sales tax is lower, but that whether and by how much borrowing falls depends on the LTV. This scheme abstracts from any effect of the decline in the sales tax on housing demand, and in so doing takes house prices as given.

We do not correct LTVs up to 100 percent, which is consistent with the fact that in that range, there are no noteworthy differences between the pre-cap and 106-cap distributions, as plotted in the top panel of Figure 1A. This feature of the data suggests that neither the imposition of the LTV limit nor the reduction in the sales tax substantially affected mortgage borrowing in that range. This could happen, for instance, if households treat any mortgage borrowing up to the house price as a way to finance the purchase of the house itself, as opposed to closing costs or home improvements.

`https://www.afm.nl/nl-nl/consumenten/veelgestelde-vragen/hypotheken-uitzondering-regels`.

[18] The LLD provides the origination month but not the day. We do not correct LTVs dated June 2011. Our not correcting LTVs for the first two weeks of that month imposes the assumption that no household registered a house purchase in the last two weeks before the reduction in sales taxes.

Recall that in the Netherlands, buyer-incurred closing costs typically amounted to about 10 or 11 percent of the house price when the sales tax was at 6 percent. Against this background, we assume that for LTVs exceeding 100 but below 110 percent, mortgage borrowing fully finances the house purchase and partially finances closing costs. In addition, we assume that households allocate any mortgage borrowing in excess of the house price equally among the components of closing costs. In that spirit, we correct LTVs exceeding 100 but below 110 by part of the decline in the sales tax. For instance, for a borrower who had an LTV of 105 with the 6 percent sales tax, we assume that mortgage borrowing covers half of the tax such that our adjustment implies that the LTV declines by half of the decline in the sales tax.

When $\widetilde{LTV}_{it} \geq 110$, we assume that mortgage borrowing covers all closing costs. In that range, we assume that relative to the uncorrected measure, the corrected LTV falls by the full extent of the decline in the tax.

Now, we explain our method for translating changes in the LTV limit into changes in the average LTV. We measure the ceteris paribus effect of a change in the limit from period $t-1$ to t as follows. First, we select a proxy for what the LTV distribution would have been in t if the cap had been at its $t-1$ level. Next, we apply the period-t cap to that distribution. We capture the effect of the change in the cap on the distribution through its effect on the cross-sectional average.

We describe two variants and apply these to Dutch LTV policy. The first is ex ante in that it only requires data from before the imposition of the cap and can therefore be applied even if an LTV limit is not actually in place. The second approach is ex post in that it also uses information that arrived after the imposition of the cap.

We first explain the ex ante approach. We start with the distribution for the last twelve months before the cap was introduced (2010M8-2011M7), after correcting for the 2011 cut in sales taxes in the way we just described. This is the tax-adjusted pre-cap distribution, plotted in blue in the top panel of Figure 2. In this distribution, the mean LTV is 93.72 percent. We use this distribution as a proxy for what the LTV distribution would have been in 2011M8-2012M12 if there had been no cap in that period either. To measure the effect of the imposition of the cap, we set all LTVs exceeding 106 in the tax-corrected pre-cap distribution to a value of 106. This yields the pink distribution in the top panel of Figure 2. The mean of the resulting distribution is 92.57 percent. Therefore, in this setting we find that the effect of the imposition of the LTV limit is to reduce the average LTV by 1.15 percentage points.

This change in the average LTV occurred on August 1, 2011. Therefore, for the purpose of measuring the annual change in the average LTV implied by the imposition of the cap, we assign an LTV of (7/12) 93.72 + (5/12) 92.57 to 2011. The cap was 106 throughout 2012, so in this context we assign a value of 92.57 to that year. Using these figures for the average LTV, we compute that the imposition of the limit implied a decline in the average LTV by 48 basis points in 2011 and by 67 basis points in 2012.

Next, we measure the effect of the subsequent reduction in the cap from 106 to 105 percent in 2013. As a measure for what the LTV distribution would have been in 2013 if the cap had still been at 106 at that time, we use the distribution that we constructed in the previous step by capping the 2010M8-2011M7 distribution at 106. We repeat this distribution, but now in blue, in the bottom panel of Figure 2. Recall that the average of this distribution is 92.57. Next, we impose a cap of 105 on this distribution by replacing all LTVs exceeding 105 by a value of 105. This yields the pink distribution in the bottom panel of Figure 2. The average of the resulting distribution is 92.35. In this context, we therefore conclude that the reduction of the cap from 106 to 105 percent implied a decline in the average LTV by 22 basis points.

We proceed analogously for subsequent one percentage point reductions in the cap through the tightening of the cap to 100 in 2018. In essence, we impose ever tighter LTV limits on the tax-adjusted 2010M8-2011M7 distribution. Table 1 reports the results. The column "μ" reports the level of the average LTVs, in percent, for the actual tax-adjusted pre-cap distribution as well as for the distributions we constructed by capping that distribution at all integer values ranging from 106 through 100. The right part of the table reports the implied annual changes in the average LTV, in percentage points.

According to the ex ante approach, the imposition and gradual reduction of the LTV limit in 2011-2018 implied a cumulative decline in the average LTV by 3.08 percentage points. We state this figure at the bottom right of Table 1 and we computed it as the sum of the annual effects tabulated in the rightmost column.

Let us now focus on the effect of the changes in the limit that occurred after the initial imposition of the cap. Recall that the decline in the cap from 106 to 105 percent implied a decline in the average LTV by 22 basis points. As the rightmost column of Table 1 documents, the ensuing one percentage point tightenings in the limit had ever larger effects. For instance, the cut in the cap from 101 to 100 percent implied a decline in the average LTV by 42 basis points. This illustrates that, for a given underlying LTV distribution, a change in the LTV limit from a lower starting point has larger effects on the average LTV. In that sense, changes in the LTV limit have a non-linear effect on the average LTV.

Figure 2: Ex Ante Approach: Imposition and Subsequent Reduction of the LTV Limit

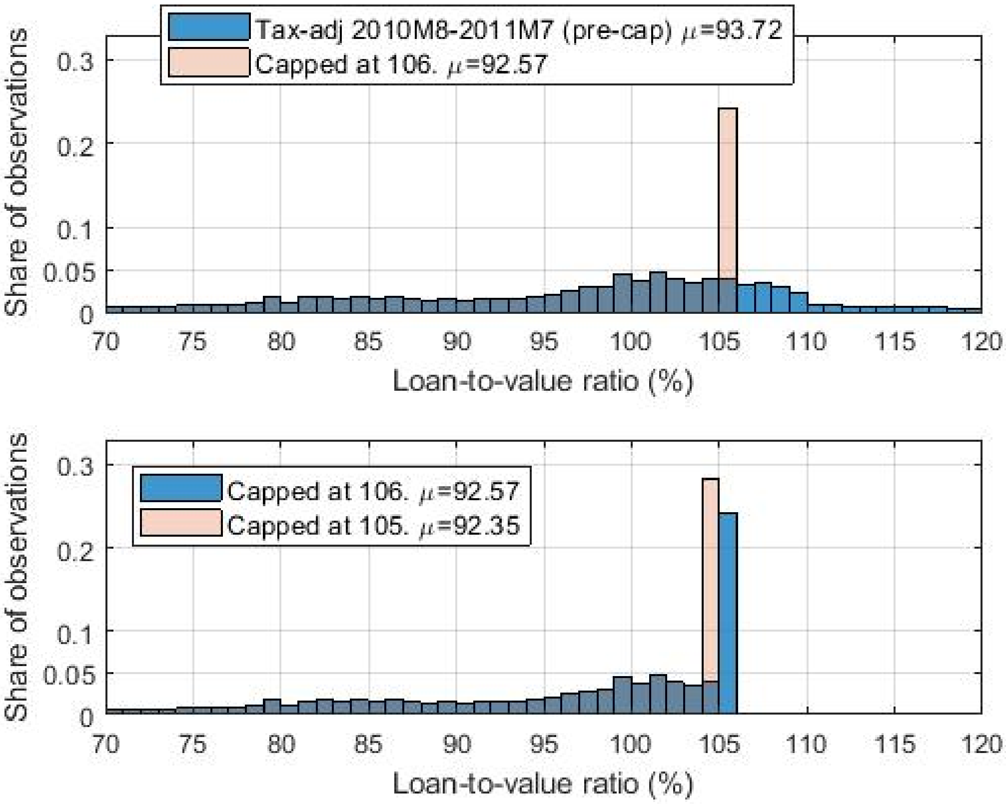

Note: This figure illustrates the ex ante approach. In the top panel, blue indicates the empirical tax-adjusted pre-cap distribution (2010M8-2011M7) while pink is that same distribution capped at 106. In the bottom panel, blue repeats the distribution which we capped at 106 from the top panel, while pink indicates the same distribution after capping at 105. For each of the plotted distributions, the legends indicate the cross-sectional mean μ . Table 1 below documents how we use these means to compute the effect of changes in the LTV cap on the average LTV.

Table 1: Ex Ante Approach: Effect of Changes in Loan-to-Value Limit on Average Loan-to-Value Ratio

Distribution		μ	Year	Change in Average LTV
Tax-adjusted pre-cap		93.72		
Capped at...	106	92.57	2011	-0.48
	/	/	2012	-0.67
	105	92.35	2013	-0.22
	104	92.09	2014	-0.26
	103	91.79	2015	-0.30
	102	91.45	2016	-0.34
	101	91.07	2017	-0.38
	100	90.65	2018	-0.42
			cumulat.	-3.08

Note: This table documents how we apply the ex ante approach to measuring the effects of the changes in the Dutch LTV limit in 2011-2018 on the average LTV. We impose ever tighter caps on the pre-cap distribution of 2010M8-2011M7, after adjusting that distribution for the 2011 change in sales taxes according to system (1). The column labeled "μ" diplays the cross-sectional average of the tax-adjusted distribution of 2010M8-2011M7 as well as the averages after capping this distribution at a level ranging from 106 through 100 as indicated in the preceding column. The rightmost column presents the implied annual changes, in percentage points, in the average LTV. We computed these from the figures in the column "μ" as the difference between the average loan-to-value ratio under the LTV limit that applied in that year and the average that applied in the previous year. Since the imposition of the cap occurred in the course of 2011, we assign its effect in part to 2011 and in part to 2012, as we explain in the body of the text. The cell labeled "cumulat." indicates the full cumulative decline in the average LTV implied by Dutch LTV cap changes in 2011-2018.

Figure 3: Ex Post Approach: Imposition and Subsequent Reduction of the LTV Limit

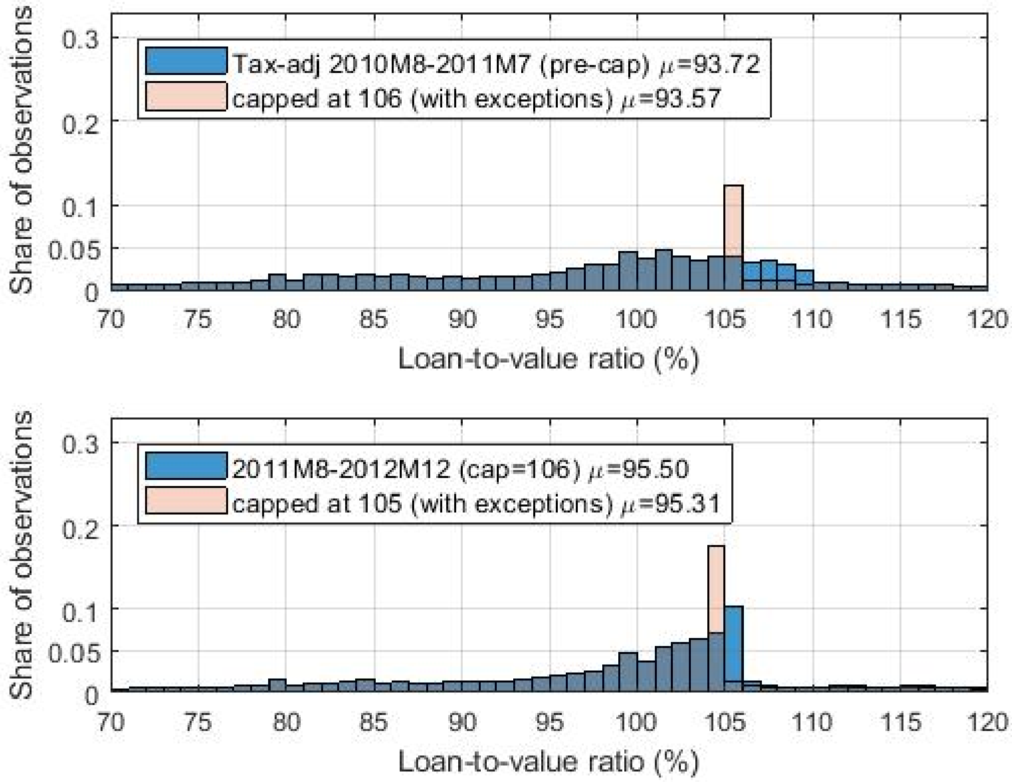

Note: This figure illustrates the ex post approach. The top panel compares the empirical tax-adjusted pre-cap distribution (2010M8-2011M7) to the same distribution which we capped at 106, while allowing for exceptions to the cap in such a way that the share of observations exceeding the cap is similar to that of the empirical 106-cap distribution. The bottom panel compares the empirical 2011M8-2012M12 distribution, when the cap was 106, to that same distribution after capping it at 105, while allowing for exceptions. The legends indicate the mean μ for each of the plotted distributions. Table 2 below documents how we use these means to compute the effect of changes in the LTV cap on the average LTV.

Table 2: Ex Post Approach: Effect of Changes in Loan-to-Value Limit on Average Loan-to-Value Ratio

Empirical distributions	μ_1		Capped at...	μ_2	Year	Change in Average LTV
Tax-adjusted pre-cap	93.72	→	106	93.57	2011	-0.07
/	/		/	/	2012	-0.09
106-cap	95.50	→	105	95.31	2013	-0.19
105-cap	93.14	→	104	92.93	2014	-0.21
104-cap	92.64	→	103	92.33	2015	-0.31
103-cap	94.44	→	102	94.06	2016	-0.38
102-cap	93.42	→	101	93.01	2017	-0.41
101-cap	88.79	→	100	88.48	2018	-0.31
					cumulat.	-1.96

Note: This table documents how we apply the ex post approach to measuring the effect on the average LTV of Dutch LTV policy in 2011-2018. The column "μ_1" displays the cross-sectional average of empirical Dutch LTV distributions ranging from the tax-adjusted pre-cap distribution of 2010M8-2011M7 to the distribution for 2017, when a cap of 101 was in place. The column "μ_2" shows the average of those distributions after imposing the indicated cap. The righmost column presents the implied changes in the average LTV, which we computed as the difference between the values for μ_2 and μ_1 in that same row. The cell labeled "cumulat." indicates the total decline in the average LTV implied by Dutch LTV policy in 2011-2018. Because the imposition of the cap occurred in the course of 2011, we assign its effect in part to 2011 and in part to 2012, as we explain in the body of the text. In the ex post approach, we account for exceptions to the cap in ways we explain in the body of the text.

To understand why this happens, notice in the bottom panel of Figure 2 that the fraction of observations in the 105-106 bin in the distribution capped at 106 (in blue) is somewhat smaller than the fraction in the 104-105 bin in the distribution capped at 105 (in pink). Throughout the gradual reduction to an LTV of 100 percent (not plotted), the size of the bar containing the limit continues to rise as the cap tightens further. This reflects the fact that as the cap tightens, it constrains a larger fraction of households. This fraction is roughly 20 percent at a limit of 105, but it is about 40 percent at a limit of 100. The larger is the fraction of households that is already constrained by a given cap, the larger is the fraction of LTVs that declines along with any decline in that cap. This in turn implies a larger effect on the average LTV.

The approach we discussed thus far was ex ante in the sense that it exclusively uses information from before the cap was introduced. In some settings, such as when a cap is not yet in place, the ex ante approach is the only option. In our application, which involves gauging the effects of Dutch LTV policy in 2011-2018, we do have more recent distributions available. In the ex post approach, we use that information.

The first way in which we use information that was not available before the imposition of the cap is in accounting for exceptions to the limit, a topic which we touched upon at the end of Section 2.1. The starting point for computing the effect of imposition remains the tax-corrected pre-cap distribution of 2010M8-2011M7. Recall that the mean of this distribution is 93.72 percent.

Now we impose a limit of 106 on this distribution while accounting for exceptions to the cap. We do so by selecting observations that are not to be moved to the cap. We do so in a way that is meant to replicate the actual share of observations that remained above the cap in the empirical 106-cap distribution.

First of all, we find that 8.12 percent of LTVs exceed 110 in the empirical tax-adjusted pre-cap distribution while 8.39 percent does in the empirical 106-cap distribution. These fractions are very similar, which suggests that the imposition of the cap had little effect on LTVs above 110.[19] With this information in hand, we assume that all observations with LTVs above 110 are exceptions to the cap that remain unaffected when we impose a limit of 106 percent on the pre-cap distribution.

Next, we focus on LTVs exceeding 106 but not exceeding 110. In the tax-adjusted pre-cap distribution, 12.03 percent of observations falls in this range while only 3.58 percent does in the empirical 106-cap distribution. With this information, we impose the limit on the pre-cap

[19] There are plausible reasons why very high LTVs are more likely to fall under the two exceptions to the limit applying to starters. Our intuition is that borrowers to whom banks lend at high LTVs tend to have a high income such that they are more likely to fall under the income-based exception to the LTV limit. Furthermore, our intuition is that very high LTVs are more likely to be partly used towards funding energy-saving home improvements, and thus are more likely to fall under that exception. This is because LTVs above 110 are bound to cover more than the house and closing costs, and therefore are used towards some other expenditure.

distribution while leaving a number of observations unaffected in such a way that the share of observations that remains in the 106-110 range is the same as that in the data, i.e. 3.58 percent. Concretely, we randomly select a fraction 3.58/12.03, i.e. 29.77 percent of observations in the 106-110 range to remain unchanged even after imposition of the cap, while setting other LTVs in that range to a value of 106 percent.

The top panel of Figure 3 illustrates how we impose the cap of 106 while accounting for exceptions. When we impose the cap while accounting for exceptions, the average LTV is 93.57 percent. This constitutes a drop by 16 basis points relative to the pre-cap distribution. In analogy with the ex ante case, we assign this effect to the two affected years such that five twelfths of it occur in 2011, which implies a drop by 7 basis points, while the remainder, a 9 basis point drop, occurs in 2012.

The effects of imposition of the cap are therefore much smaller than in the ex ante approach. We return to this issue at the end of this section.

The second way in which we use information that only became available ex post is that we use the most recent prior empirical distribution as a proxy for what the distribution would have been if the cap had been at its prior level. For instance, to gauge the effect of the cut in the LTV limit from 106 to 105 percent in 2013, we use the empirical 106-cap distribution of 2011M8-2012M12 as a measure for what the LTV distribution would have been in 2013 if the limit had still been at 106. In so doing, we use the most recent distribution prior to the cap of 105. This reflects our intuition that more recent distributions are in general a better proxy for the underlying current intentions of lenders and borrowers. The ex post approach takes into account the fact that the underlying LTV distribution can have changed for various reasons other than the current change in the cap, including but not limited to changes in expected household income, changes in the dispersion of income expectations and changes in other prudential policy instruments such as capital requirements on banks. Another point relates to the effect of past changes in LTV caps. While we discussed in Section 2.1 that we see little evidence that LTV limits affected LTVs below the cap, that point pertains to the effect on impact. It might still be the case that limits affect LTVs below the cap with a lag, such that it is still relevant to note that the ex post approach takes into account any such indirect effects of past changes in the cap on LTVs below the cap.

The average of the empirical 106-cap distribution is 95.50 percent. We impose a cap of 105 percent on this distribution by setting all LTVs above 105 but not exceeding 106 to a value of 105. In doing so, we abstract from any exceptions in that particular bin. We treat LTVs above 106 in the empirical 106-cap distribution as exceptions to the limit, that is, we do not move them to 105. Still, we assume that a one percentage point reduction in the cap causes these observations to shift downward by one percentage point. We therefore assume that excepted LTVs are still set relative to the current limit.

When we impose a cap of 105 on the empirical 106-cap distribution in this fashion, the cross-sectional average falls to 95.31. This implies that the one percentage point reduction in the cap to 105 implied a reduction in the average LTV by 19 basis points. The bottom panel of Figure 3 illustrates how we impose the cap of 105 while accounting for exceptions.

We proceed analogously through the lowering of the cap from 101 percent to 100 percent in January 2018. In this last case, we impose the 100 percent cap, as it applied in 2018, on the empirical LTV distribution of 2017.

Table 2 details the ex post approach as applied to Dutch LTV policy in 2011-2018. The column labeled "μ_1" contains the averages in successive empirical distributions, while the column labeled "μ_2" contains the averages after we imposed the cap on the empirical distribution in the same row, while accounting for exceptions. The rightmost column reports the implied change in the average LTV, computed as the difference between the values for μ_2 and μ_1 in the same row.

The cumulative effect on the average LTV is 1.96 percent. This is about one percentage point less than in the ex ante approach. Comparing the rightmost columns of Tables 1 and 2, we find that this difference is largely due to the difference in the measured effects in 2011 and 2012, i.e. the effects of the imposition of the cap. Recall that in both cases, we took the same pre-cap distribution as a starting point for the imposition of the cap. In addition, we imposed the same limit of 106. As for the imposition of the cap, the only difference was that we allowed for exceptions in the ex post approach while we did not in the ex ante approach. Therefore, the difference between the ex ante and ex post approach in terms of the effects in 2011 and 2012 can only be due to the different treatment of exceptions.

When we sum the changes in the average LTV in 2013 through 2018, all of which are implied by changes in an existing limit, we find a cumulative decline by 1.93 percentage points in the ex ante approach and by 1.80 percentage points in the ex post approach. These are similar, which suggests that in our application, one could capture the cumulative effect of changes in the limit quite well by using only the pre-cap distribution.

The rightmost column of Table 2 shows that the effects of changes in the limit still tend to increase as the cap tightens, but also that this relation is not monotonic. This is because in this setting, we apply the respective LTV limits to different distributions. For instance, the empirical 101-cap distribution lies further to the left than the previous distribution. This tends to imply that the limit is binding for a smaller share of households, which tends to lower the effect of the cap on the average.

3 Time Series Data

In this section, we discuss the time series which we use, with an emphasis on our measure for the average loan-to-value ratio. Figure 4 plots the series. All are at an annual frequency, for reasons related to the sample size of the survey using which we computed the average LTV.

The upper left panel shows our measure for the average loan-to-value ratio, which we discuss in a minute. The upper right diagram plots the real mortgage interest rate in percentage terms. We compute the real rate r_t as $i_t - (1/2) \sum_{j=0}^{1} \pi_{t-j}$, that is to say, as the difference between the nominal mortgage rate i_t and a two-year moving average of inflation π_t in the GDP deflator.[20] We use this moving average as a proxy for inflation expectations.

The lower left panel shows real GDP in billions of 2010 euros.[21] The lower right panel shows a constant-quality index for the relative price of housing. It is based on a nominal house price index known as the "Prijs Bestaande Koopwoningen" (price of the existing housing stock).[22] We compute the relative house price by deflating the nominal index by the GDP deflator.

The remainder of this section discusses our measure for the average loan-to-value ratio. We measure the LTV for first-time buyers at the time of purchase. One reason for focusing on respondents' first housing purchase is that unlike what is the case for subsequent purchases, it is not affected by the extent of any previous housing equity accumulation, whether it be positive or negative equity. In their studies involving a survey-based LTV measure for the United States, Duca, Muellbauer and Murphy (2011, 2012) similarly measure the LTV for first-time buyers.

[20] The nominal mortgage rate series which we use is the standard series used by De Nederlandsche Bank. From 2003Q1 on, it is the weighted average of the residential property mortgage rates for different maturities. Before that date, it is the rate for five-year mortgages.

[21] Source: Statistics Netherlands and De Nederlandsche Bank.

[22] The Prijs Bestaande Koopwoningen (PBK) is from Statistics Netherlands and the Netherlands' Cadastre, Land Registry and Mapping Agency. Constant-quality data for this series are only available from 1995 on. For earlier years, we extrapolate the PBK backwards using the growth rate of the constant-quality index from Bussel, Kerkhoffs and Mahieu (1996).

Figure 4: Macro Series

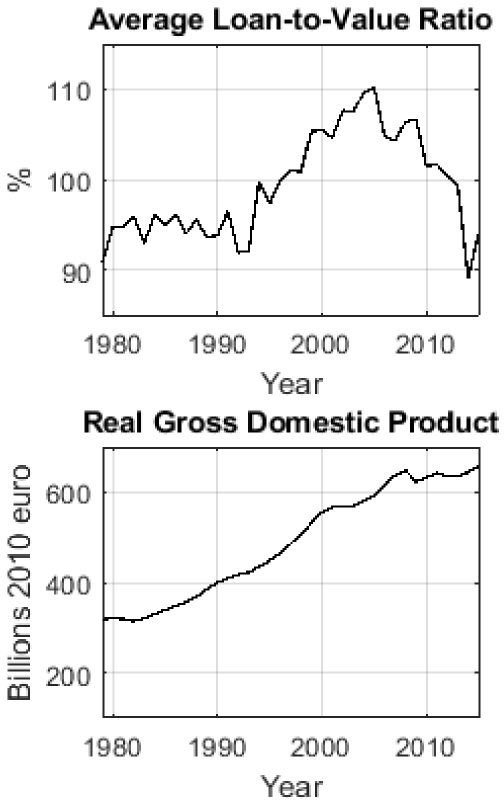

Note: This figure plots annual time series (1979-2015) for the Netherlands. Our survey-based average loan-to-value ratio is in percent. The real mortgage interest rate (in percent) is the nominal mortgage rate minus the two-year moving average of inflation in the GDP deflator. Real Gross Domestic Product is in billions of 2010 euro. The relative house price is a constant-quality index for the nominal price of housing based on the "Prijs Bestaande Koopwoningen" (price of the existing housing stock) deflated by the GDP deflator. Sources: CentERdata; De Nederlandsche Bank; Statistics Netherlands; the Netherlands' Cadastre, Land Registry and Mapping Agency; Bussel, Kerkhoffs, and Mahieu (1996) and our own calculations.

To measure the LTV, we designed a household survey that we fielded in the CentERpanel as well as in the LISS panel. Both of these household panels are administered by CentERdata at Tilburg University.[23] Both panels are designed to be representative of the Dutch-speaking population of the Netherlands.

CentERdata fielded our surveys in January 2016. We selected all individuals in the CentER and LISS panels of age eighteen and older, irrespective of their position in the household to which

[23] See www.centerdata.nl and www.lissdata.nl for a complete description of the two panels.

they belong. We fielded the same survey in both panels. Questions and responses were in Dutch.

In the CentERpanel, 2,368 individuals responded, which implies a response rate of 92.72 percent. In the LISS panel, 5,575 panel members responded, at a response rate of 83.28 percent. These figures for the number of respondents and response rates are including partially filled in questionnaires.

Given that the population of the Netherlands amounts to a rounded seventeen million, we consider it highly unlikely that there is non-trivial overlap between the two samples. We therefore merge the two samples without attempting to control for any overlap. This yields a total of 7,943 respondents.

Our questionnaire pertains to respondents' first purchase of residential property. The appendix provides an English translation of relevant questions. In this paper, we exclusively use responses from individuals who report ever having bought a house or apartment.

In an effort to enhance precision of the answers, the introduction to the questionnaire asks respondents to fetch any documents which contain the purchase date of their first property, the purchase price and the mortgage loan amount.

We ask any respondent who ever bought a property whether he or she bought his or her first house or apartment alone, with the current partner, with a former partner, or in any other association. If the property was purchased with a current or former partner, we ask whether it was also the first home purchase for that partner. We only include purchases which were carried out alone, with the current partner or a former partner and where the single buyer or both buyers were first-time buyers. We avoid double counting by including at most one observation when the first home purchase is carried out with the current partner.[24]

[24] We deal with doubles as follows. Potential doubles only occur when a respondent reports buying his or her first property with the current partner. We assume that two such respondents only report on the same purchase if both belong to a category which includes the household head and his or her married or unmarried partner, or if neither

Next, we ask in which year and month the buyer(s) signed the purchase agreement. Throughout this paper, we restrict attention to annual data because when we used information on the month of purchase to construct a quarterly LTV measure, the quarterly sample size turned out to be very small.

Subsequently, we ask for the purchase price of the first home. We ask respondents to enter the amount agreed between buyer and seller excluding buyer-incurred closing costs.

Immediately after the purchase price question, we ask whether respondents took out a mortgage loan in connection with the purchase of their first home. Consistent with the fact that the LLD data from Section 2 contain only individuals that took out a mortgage loan, we omit the small fraction of survey respondents who bought their first home without a mortgage loan.

Next, we ask a question providing a choice out of thirteen possible responses regarding the level of the mortgage loan amount relative to the house price. The options range from the total mortgage loan amount being "more than 50 percent below the purchase price" over "about equal to the purchase price" to "more than 30 percent above the purchase price". We ask respondents to select their answer such that it reflects the sum of all loan amounts in case the home purchase was financed by multiple mortgage loans.

We positioned and designed this question so as to place a low burden on respondents. We ask it immediately after the house price question to make it easier for respondents to compare the loan amount to that price. The question does not require that respondents have a precise recollection of the house price or the loan amount. Moreover, they do not need to fill in a particular percentage, but rather can choose from a range of cases with approximate percentages.

of the two belongs to that category. To drop one of each pair of doubles, we adopt the following two steps. First, if our main cleaned LTV measure (which we refer to below as the "computed LTV after cross-checking") is available for one respondent but not for the other, we omit the latter's responses. For remaining doubles, we randomly select the response of one partner with each having a probability of 50 percent of being selected.

We assign a particular value for the LTV to every of the thirteen responses. For responses indicating that the total loan amount was “more than 30% above the purchase price”, we assume an LTV of 140 percent, while for the case “more than 50% below the purchase price”, we assume that the LTV was 25 percent. In all other cases, we assume that the LTV equals the approximate value that is indicated in the question. For instance, for a loan amount “about 10% above the purchase price”, we assume an LTV of 110 percent.

We call the LTV which we thus constructed the “directly reported LTV”.

We also gauge the LTV in another way. If a respondent answered the question on the purchase price and reports taking out a mortgage, we ask for the total mortgage loan amount associated with the purchase of the first home. As before, we ask to sum the amount of any multiple mortgages. Implicitly, both the loan amount and the house price are in nominal terms. In this case, we compute the LTV as the ratio of the reported loan amount to the reported house price, expressed in percentages. We call this the “computed LTV”.

Our intuition is that for respondents who remember the house price and the loan amount well or looked up that information as we requested, the computed LTV is closer to the truth than the directly reported LTV. This is because the latter rests on an approximate figure while the former does not. On the other hand, our sense is that for respondents who do not remember the relevant information well, large discrepancies from the truth are more likely with the computed LTV than with the directly reported LTV. Amongst others, this is because the latter is designed to be easier to answer and because the coding scheme places bounds on measured LTVs.

Aiming to combine the benefits of both measures, our baseline measure primarily relies on the computed LTV but uses the directly reported LTV as a cross-check. In particular, we drop observations for the computed LTV where the computed LTV differs by more than 25 percentage points from the directly reported LTV.[25] We call this the "computed LTV after cross-checking".

Furthermore, we omit observations for which the computed LTV exceeds 200 percent. There are no reported negative loan amounts.

In the ensuing sections, we use data ranging from 1979 to 2015. Within this time span, there are 2,238 observations in our cleaned sample for the computed LTV after cross-checking. This corresponds to an average of 60.49 observations per year.

In Figure 4 above, the top left panel displays the cross-sectional average of the computed LTV after cross-checking. On average in 1979-2015, this measure for the LTV for first-time buyers is 99.23 percent. It hovered between 90 and 97 percent in the 1980s and early 1990s. It then gradually increased from a trough of 91.96 percent in 1992 to its maximum of 110.17 percent in 2005. It has since declined, attaining 94.02 percent in 2015.

These figures reflect the fact that loan-to-value ratios in the Netherlands are high in comparison with those in other economies.[26] One of the plausible reasons for this is the fact that mortgage interest payments can be deducted from taxable income at a particularly high rate.[27] Another factor which likely plays a role is the existence of a government-backed mortgage default insurance scheme for dwellings with a purchase price below a certain threshold, called the Nationale Hypotheek Garantie (NHG). This fund currently covers about one third of outstanding mortgages in the Netherlands.[28]

[25] We treat the two outer cases of the directly reported LTV as follows. For the case "more than 30% above the purchase price", we dropped computed LTVs below 115 percent. For the case "more than 50% below the purchase price", we dropped computed LTVs above 65 percent.

[26] Table 3 of Catte, Girouard, Price, and André (2004) indicates that the Netherlands has the single highest value for the "typical LTV" in a set of 18 developed economies. Table 1 of Calza, Monacelli, and Stracca (2013) indicates that among 19 industrialized economies, the typical LTV is highest in the Netherlands, with the nuance that the typical Dutch LTV equals the *upper end* of the range for the typical LTV in the United Kingdom.

[27] According to European Commission (2014), mortgage interest tax relief lowers the user cost of housing by more in the Netherlands than in any of the other 25 European Union member states included in its Graph 3.5. The maximum tax rate applicable for the deduction of mortgage interest expenditure was 52 percent as of 2013 and is since being gradually lowered.

[28] As of 2017, NHG covered 1.3 out of 4.3 million owner-occupied dwellings in the Netherlands, where the latter figure includes dwellings on which no mortgage is outstanding. NHG has been operational since January 1, 1995.

4 Macroeconomic Effects of Changes in the Loan-to-Value Limit

This section estimates the effects of Dutch LTV policy actions in 2011-2018 on house prices and real activity.

We use the four series plotted in Figure 4. Table 3 reports Augmented Dickey-Fuller (ADF) t-statistics along with the associated p-values and the lag length in the testing regressions, where we used the Schwarz Information Criterion (SIC) to select lag length. The top part of the table reports results for the levels of the average loan-to-value ratio (LTV_t), log relative house price (q_t), log real GDP (y_t) and the real mortgage rate (r_t). The top left part of the table pertains to ADF regressions including an intercept as the only exogenous variable while the top right part pertains to regressions including an intercept and a linear time trend. The bottom part reports results for the same variables in first differences, with the left part pertaining to regressions without intercept or trend and the right part to regressions with an intercept only. All test regressions reported in Table 3 are for 1981-2015, with additional years used for differencing and lags.

Whether we include a deterministic trend or not, we do not reject the null hypothesis of a unit root in any of the four level series, not even at the 10 percent level. For the differenced variables, we reject the null hypothesis of a unit root in all reported cases at the 5 percent level or better, except when we include an intercept in the regression for relative house price inflation. In that case, the ADF t-statistic is -2.51 and the p-value is 0.12.

From 1956 until that time, there was a municipal mortgage guarantee scheme. See the Dutch-language Staatsblad 1995, No. 28 for information on the essence of both schemes.

Table 3: Augmented Dickey-Fuller Stationarity Tests

	k			k and t		
	ADF t-stat	p-value	lag length	ADF t-stat	p-value	lag length
LTV_t	-1.68	0.43	0	-1.36	0.86	0
q_t	-1.56	0.49	1	-2.60	0.28	3
y_t	-0.95	0.76	1	-1.65	0.75	1
r_t	-0.81	0.81	0	-1.98	0.59	0

	neither k nor t			k		
	ADF t-stat	p-value	lag length	ADF t-stat	p-value	lag length
ΔLTV_t	-7.79**	0.00	0	-7.68**	0.00	0
Δq_t	-2.42*	0.02	0	-2.51	0.12	0
Δy_t	-1.96*	0.05	0	-3.17*	0.03	0
Δr_t	-5.75**	0.00	0	-5.77**	0.00	0

The top half of this table reports the Augmented Dickey-Fuller (ADF) t-statistic, its associated p-value, and the lag length for ADF regressions in the levels (top half of the table) and first differences (bottom half) of the loan-to-value ratio (LTV_t), log relative house price (q_t), log real GDP (y_t) and the real mortgage rate (r_t). We used the Schwarz Information Criterion to select lag length. The series are those from Figure 4, but after dividing the LTV and the mortgage rate by 100, after dividing house prices by 100 and then taking logs, and after taking logs of real GDP. The sample period is 1981-2015, with earlier years used for differencing and lags. "k and t" indicates that we include both a constant and a linear deterministic trend in the ADF regression; "k" indicates that we include a constant but no trend; "neither k nor t" means we do not include either. The ADF tests do not reject a unit root in either of the levels variables. At the five percent level, we reject a unit root in the differenced variables in all cases except the regression for the differenced log relative house price with a constant. * indicates statistical significance at the 5 percent level; ** indicates significance at the 1 percent level.

When we use the Akaike Information Criterion (results not reported here), we again cannot reject a unit root in any of the levels variables, whether we include a deterministic trend or not. In that case, we do reject the null of a unit root in each of the four differenced variables, including relative house prices, at the five percent level or better, whether we include an intercept or not.

With this evidence in mind, we treat all four series as I(1).

Next, we implement Johansen (1988) cointegration rank tests. In both the cointegration relation and the equations of the Vector Error Correction Model (VECM), we include intercepts but no deterministic trend. We perform these tests using a VECM with one lag, which is the value selected by the SIC. The results are in Table 4. Both L-Max and Trace tests reject the null hypothesis of no cointegration at the 1 percent level. Neither the L-Max nor the Trace test indicate the presence of multiple cointegration relationships.

Based on this evidence, we estimate a VECM with a single cointegrating vector using the Johansen maximum likelihood estimator. The estimation period is 1981-2015, with additional years used for differencing and lags. We normalize the cointegration equation such that the coefficient on house prices is one. It is therefore as follows:

$$q_t - k - \omega_y \ y_t - \omega_r \ r_t - \omega_{ltv} \ LTV_t = \eta_t \tag{2}$$

where q_t is the log relative house price, y_t is log real GDP, r_t is the real mortgage rate, LTV_t is the average LTV, k is a constant and η_t is the cointegration residual. ω_y, ω_r and ω_{ltv} are the long-run coefficients.

Based on the multivariate SIC, we choose a VECM with 1 lag. Our VECM is as follows:

$$\Delta X_t = A + \Delta X_{t-1} + \alpha\beta' X_{t-1} + \varepsilon_t \tag{3}$$

where $X_t = (LTV_t, \ q_t, \ y_t, \ r_t)'$ contains the data, A is a vector of constants, β is a vector with the estimated long-run coefficients, such that the error correction term $\beta' X_{t-1} = \eta_{t-1}$, α is a vector of coefficients capturing whether and how strongly each variable corrects for any deviation from long-run equilibrium, and ε_t is a vector of reduced-form short-run residuals.

Table 4: Johansen Cointegration Rank Tests

Test	H_0	H_A	$\lambda_{\max}$	p-value
L-max test	r=0	r=1	37.88**	0.00
	r=1	r=2	17.69	0.14
	r=2	r=3	4.79	0.77
	r=3	r=4	0.39	0.53
Trace test	H_0	H_A	λ_{trace}	
	r=0	r> 0	60.76**	0.00
	r≤1	r> 1	22.88	0.25
	r≤2	r> 2	5.18	0.79
	r≤3	r> 3	0.39	0.53

Note: This table presents Johansen L-max and Trace test statistics for the number of cointegrating relations in the system containing the four variables from Table 3. The tests are based on a VECM specification with one lag and with an intercept but no deterministic trend in both the short-run equations and in the long-run relation. The Schwarz Information Criterion suggests using one lag. At the 5 percent level, both the L-max and Trace test reject the null hypothesis of no cointegration and do not indicate the presence of multiple cointegration relations. ** indicates significance at 1% level; * indicates significance at the 5% level.

We use the estimated VECM to examine the macro effects of the introduction and gradual reduction of the Dutch LTV limit. Recall that in Section 2.2, we translated the changes in the Dutch LTV limit that occurred in 2011-2018 into changes in the average LTV. We now feed these changes in the average LTV into the VECM as a succession of eight successive annual LTV impulses. We take these impulses from the rightmost columns of Table 1 and 2 for the ex ante and ex post approach, respectively.

In the baseline, we use the following Cholesky ordering: [LTV, house prices, GDP, mortgage rate]. By ordering the LTV first, we do not impose any zero restrictions on the contemporaneous responses of the other variables to LTV shocks. We impose the restriction that house prices do not instantaneously react to GDP and the mortgage rate based on the fact that house prices are typically persistent such that any instantaneous response to other variables is likely to be small. We impose the restriction that GDP responds to mortgage rates with a lag. The ordering does not impose any restrictions on the instantaneous responses of mortgage rates, based on our intuition that banks readily adjust mortgage rates for new loans in response to economic developments.

The blue lines in Figure 5 present the accumulated impulse responses for the ex post approach, while the red lines represent the ex ante approach. The horizontal axis indicates years since the initial imposition of the cap in 2011. Since we use a succession of shocks for years between 2011 and 2018, the last shock occurs when $t = 7$ on the horizontal axis.

The impulse responses for the ex post approach imply that twenty years after the first shock, relative house prices are 4.84 percent lower, and real GDP is 1.15 percent lower, than would have been the case if there had not been an LTV limit throughout. The bulk of the responses has occurred 10 years after the first shock, with house price and GDP responses amounting to 4.18 percent and 1.05 percent, respectively. The twenty-year response of the real mortgage rate is positive but very small.

The red lines in Figure 5 show that the responses are larger when we use the ex ante approach. This occurs because the cumulative LTV impulse is larger in this case. As we discussed in Section 2, the ex post approach is the better measure when it is feasible, so the blue lines represent our best estimates of the effect of variation in the Dutch LTV cap in 2011-2018.

As we discussed in Section 2, the effect of changes in the limit depends on the relative position of the limit and the underlying LTV distribution. The limit has not fallen below 100 percent, which is high from an international perspective. Still, according to our estimates it had a non-trivial effect on house prices, which we interpret as resulting from the fact that Dutch households tend to borrow at high LTVs such that the limit was binding for a substantial fraction of borrowers.

Figure 5: Impulse Responses for Dutch Loan-to-Value Policy
from 2011 Onwards

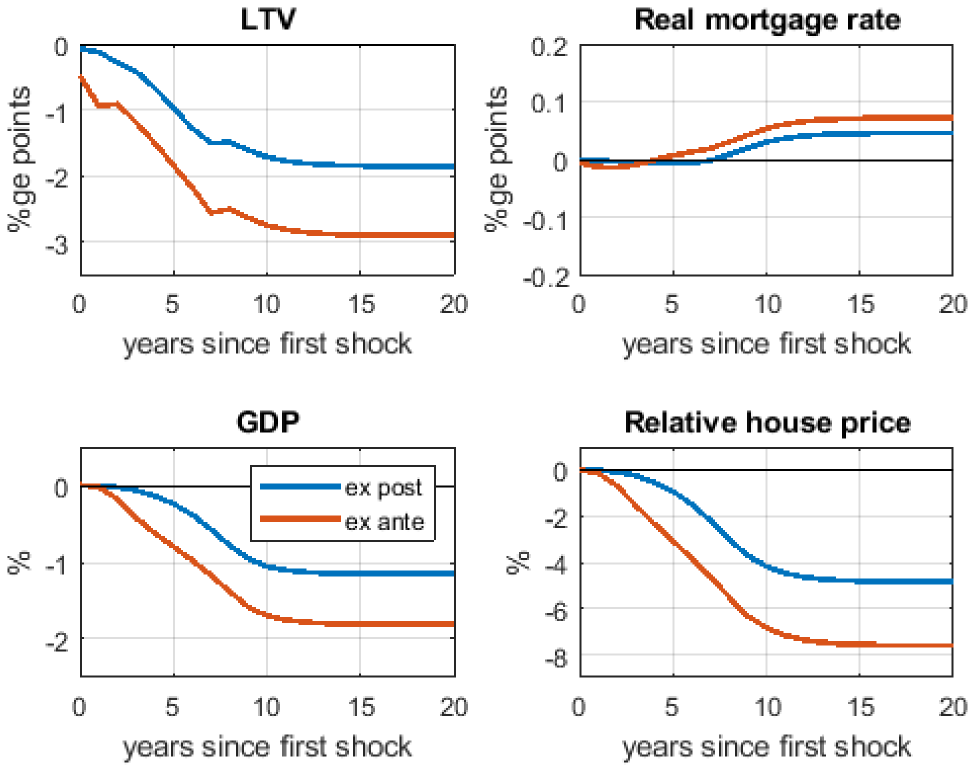

Note: This figure shows the impulse responses to a succession of eight annual impulses to the average LTV implied by the imposition of an LTV cap in the Netherlands at 106 percent in 2011 and its subsequent gradual reduction, eventually reaching 100 percent in 2018. The impulses are from the rightmost columns of Tables 1 and 2, which correspond to the ex ante and the ex post approach, respectively. We use a VECM(1) with Cholesky ordering [LTV, house prices, GDP, mortgage rate]. Blue lines reflect the ex post approach and red lines reflect the ex ante approach. Section 2.2 explains both procedures. The horizontal axis indicates the number of years since the first shock in 2011. LTV and mortgage rate responses are in percentage point deviations from the path that would occur in the absence of shocks while GDP and house price responses are in percent deviations.

As for available estimates on the effects of LTV limits in the Netherlands, CPB Netherlands Bureau for Economic Policy Analysis (2015, 2017) and De Nederlandsche Bank (2015a) examine the effects of a hypothetical reduction in the LTV limit from 100 to 90 percent. The analysis in the present paper evolved from our own contribution to De Nederlandsche Bank (2015a), which is the model with an “explicit LTV” series described in Appendix D of that study. All other procedures in the three studies we just mentioned involve imputing the decline in the number of housing transactions resulting from the decline in the LTV limit. The results from these procedures suggest that a decline in the LTV limit from 100 to 90 percent implies that in the long run, nominal house prices are lower than they would have been if the limit had stayed at 100, with the size of the long-run house price response ranging from 1 to 5 percent. In the present section we show effects for a different policy scenario, but in Section 6 below we show the long-run effects on house prices under various scenarios, including a reduction in the LTV limit from 100 to 90 percent. The relevant entry in Table 5 reveals that according to our estimates, such a reduction would imply that in the long run, relative house prices would be 14.27 percent lower than compared to a situation where the limit stays at 100 percent.

This finding suggests that our method implies larger effects on house prices than the above-mentioned procedures based on imputing changes in housing transactions. Still, a number of caveats are in order. Firstly, we do not take anticipation effects into account. To the extent that households carried out high-LTV purchases in the pre-cap period because they anticipated that their borrowing would be constrained by the cap if they bought the house later, that would tend to imply an overestimate of the effect of the cap. Secondly, this estimate is formed using the ex ante approach. As such, it does not account for exceptions to the limit, which tends to imply overestimation. Thirdly, the estimates we discussed to this point do not control for the loan-to-income ratio (LTI), which also plausibly implies overestimation.

Note that we do not assess whether loan-to-value policy has been welfare enhancing and that we do not assess what would be an optimal LTV limit. For instance, we do not assess whether and to which extent the introduction and lowering of the LTV limit has enhanced financial stability. See CPB Netherlands Bureau for Economic Policy Analysis (2015, 2017) and De Nederlandsche

Bank (2015a) for discussion and analysis of a more comprehensive set of plausible effects of LTV policy.

The remainder of this section discusses robustness. The results are robust to using all alternative Cholesky orderings and to using a real mortgage interest rate computed as the nominal rate minus current inflation. The results are robust to using the computed LTV without cross-checking as well as to using the directly reported LTV after cross-checking. When we use the directly reported LTV without cross-checking, the twenty-year response of relative house prices using the ex post approach is 3.81 percent, while that in real GDP is 0.72 percent. Recall from Section 3 that we implemented cross-checking to enhance accuracy of measurement.

Next, we check for robustness with respect to using an alternative LTV measure which equals our baseline measure through 2010 but for which we extrapolate more recent years using the cross-sectional average of the LTV from the Loan Level Data (LLD).[29] Our reasons for doing so are that the number of observations on which our baseline series is based is lowest in 2010-2015 and that the ten percentage point drop in 2014 is by far the largest one-year change in the series. The levels of the baseline and extrapolated series are similar in 2011-2013 but the extrapolated LTV drops by only half a percentage point in 2014. Using this alternative LTV series and the ex post approach, we find that the imposition of the LTV cap at 106 percent and its gradual tightening to 100 percent have somewhat larger effects than in the baseline. The cumulative twenty-year response in relative house prices is 6.65 percent and that in real GDP is 1.40 percent.

Finally, we check for robustness with respect to including the loan-to-income ratio (LTI) as a fifth variable in the VECM. We measure the LTI for first-time buyers using the same survey we used

[29] Denoting LTV_t as the cross-sectional average of our survey-based LTV in year t and $\widehat{LTV}_t$ as the annual cross-sectional average of the LTVs from the Loan Level Data (LLD), we compute the extrapolated LTV for all t from 2011 through 2015 by scaling up the LLD-based LTV as follows: $\overleftrightarrow{LTV}_t = \widehat{LTV}_t \ (LTV_{2010} \ / \widehat{LTV}_{2010})$.

for computing the LTV.[30] The LTI series is more choppy than that of the LTV, which we believe reflects both the fact that we received substantially fewer responses for the LTI and the likelihood that respondents who did report the question about their past income (or even their partner's past income) found it challenging to do so with precision. To account for at least the former of those two concerns, we proxy the LTI by a three-year backward-looking moving average. In particular, $\widetilde{LTI}_t = (1/\sum_{\tau=0}^{2} N_{t-\tau}) \sum_{\tau=0}^{2} N_{t-\tau}\ LTI_{t-\tau}$, where $LTI_{t-\tau}$ is the cross-sectional average of the LTI in year $t-\tau$ before smoothing and $\widetilde{LTI}_t$ is the smoothed LTI. Note that in computing the three-year moving average, we weigh by the number of observations $N_{t-\tau}$ in that year.

In computing impulse responses, we order the LTI second in the Cholesky ordering while we preserve the baseline relative ordering of the four other variables. In all other respects, we specify the VECM as in the baseline.

The LTI is insignificant in the long-run relation but is significant at the 5 percent level in the VECM house price equation. In the ex post approach, the twenty-year house price and GDP responses are, respectively, 3.32 and 0.74 percent. This suggests that not controlling for the LTI in the baseline model tends to cause the model to somewhat overstate the connection between the LTV and the macro-economy, with the above-mentioned caveat that we had to smooth the LTI and households likely found it challenging to answer the question about their past income with precision.

[30] Our approach to measuring the LTI is similar to that which we used for measuring the LTV. To capture the directly reported LTI, we ask respondents how their mortgage loan amount compared to the gross annual income that the lender took into account (which, depending on the case, is the respondent's own income; the partner's income; or both). Options range from "0.5 times" through "8 times" annual income in increments of 0.5, in addition to "more than 8 times" annual income. We assign an LTI of 10 to the latter case. In other cases, we use the approximate LTI that is mentioned in the survey answer. To capture the computed LTI, we ask respondents for the relevant annual income at the time of purchase of the first house. For more detail, see questions v12 through v14 in the appendix. In analogy with our approach to measuring the LTV, we construct a computed LTI after cross-checking with the directly reported LTI. We cross-check by dropping observations for which the computed LTI exceeded 1.5 times the directly reported LTI or fell below 0.5 times the directly reported LTI. However, if the directly reported LTI is 0.5, we drop the computed LTI only if it exceeds 1 or is below 0. In computing the cross-sectional average of the LTI, we take the unweighted average of all LTIs in a year.

5 The Degree of Heterogeneity Matters

In the present section, we consider the shape of the underlying LTV distribution, i.e. the distribution that would obtain if there had been no limit. In particular, we focus on its spread. We show by means of stylized scenarios that a change in the degree of heterogeneity that preserves the mean, and is as such only visible at the micro level, can substantially alter the macroeconomic effects of changes in LTV limits.

There are various reasons why heterogeneity in underlying LTVs could change over time. For instance, the loan-to-value distribution could become narrower if banks were to move from a hypothetical situation where they charged the same interest rate on all mortgage loans to a situation where the interest rate is an increasing function of the LTV on the loan. This would happen if high-LTV borrowers would respond by borrowing less and low-LTV borrowers by borrowing more. As another example, our intuition is that in a model with transitory household-specific income shocks, a rise in the variance of such shocks would widen the LTV distribution. This would happen if households that temporarily have a low income wish to borrow more while current high-income households borrow less.

To gauge the effect of the degree of heterogeneity, we generate distributions that differ in terms of their standard deviation while being essentially identical in terms of other characteristics. To this effect, we generate distributions within the Pearson system. This is a family of distributions within which every combination of the first four moments corresponds to a unique distribution. Within this system, we consider scenarios that differ in terms of the standard deviation of the distribution while having essentially the same mean, skewness and kurtosis.

We start out with the Dutch pre-cap distribution of 2010M8-2011M7 after adjusting for the 2011 decrease in sales taxes according to system (1). This distribution has mean 93.72, standard deviation 17.09, skewness -1.16 and kurtosis 5.97, where the latter is defined such that a normal distribution has a kurtosis of 3.

So as to allow for analysis within the Pearson system, we then generate a Pearson family distribution which essentially has the same first four moments as the Dutch pre-cap distribution. We refer to this as the "Pearson proxy" of the Dutch distribution. In a final step, we will alter the standard deviations of this distribution so as to construct different scenarios as to the degree of heterogeneity.

In constructing the proxy of the Dutch distribution, we account for the fact that the support of the empirical pre-cap distribution after cleaning is $(0, 200]$. This is non-trivial because the Pearson distributions are themselves non-truncated. We search across untruncated Pearson family distributions for a distribution in which, after truncation to $(0, 200]$, the first four moments are close to those of the Dutch pre-cap distribution. In particular, we use `fminsearch` to minimize the sum of the absolute deviations of the first four moments from their empirical values. Every evaluation of this function of the moments is based on a truncated Pearson distribution with ten million observations.

Minimization yields a Pearson family distribution which, before truncation, has mean 93.65, standard deviation 17.29, skewness -1.30 and kurtosis 5.69.

Next, we generate two alternative distributions which, before truncation, have exactly these values for the mean, skewness and kurtosis, but which have a standard deviation that is either 25 percent lower or 25 percent higher than 17.29. In the former case, which we call the "low standard deviation case", the standard deviation in the distribution before truncation is 12.97. In the latter, the "high standard deviation case", the standard deviation is 21.61.

We generate 10 million observations from each of the above three distributions (Pearson proxy, high and low standard deviation cases), and truncate them to the support $(0, 200]$.

In the top panel of Figure 6, blue bars represent the actual Dutch tax-adjusted pre-cap distribution while red bars indicate the Pearson family proxy. As before, we plot one percentage point bins of the LTV. The vertical axis indicates the share of observations in each bin.

In the lower panel of Figure 6, red bars repeat the Pearson proxy from the top panel. Blue and yellow bars indicate the low and high standard deviation case, respectively.

For each of the three distributions in the bottom panel of Figure 6, we examine the effects of the policy scenario that occurred in the Netherlands, i.e. the imposition of an LTV limit at 106 percent and its subsequent reduction to 100 percent in annual steps of one percent. We apply the ex ante approach from Section 2.2 on each of these three distributions so as to translate the changes in the limit into changes in the average LTV. We use the ex ante approach because we only have one distribution for each scenario. As before, in using the ex ante approach we do not allow for any exceptions to the limit.

We then feed these sets of impulses to a VECM with the same specifications and estimated parameters as in Section 4, applying the same baseline Cholesky ordering.

Figure 7 plots the impulse responses. Again, blue indicates the low standard deviation case, red the Pearson family proxy for the Dutch pre-cap distribution, and yellow the high standard deviation case. With the proxy for the Dutch distribution, relative house prices are 8.46 percent lower twenty years after the initial impulse than they would have been without any LTV impulse, while real GDP is 2.01 percent lower than it would have been. Note that these responses are larger than those for the ex post approach on which we focused in Section 4, but similar to those for the ex ante approach with the actual Dutch distribution.

Figure 6: Pearson Approximation and Varying the Degree of Heterogeneity

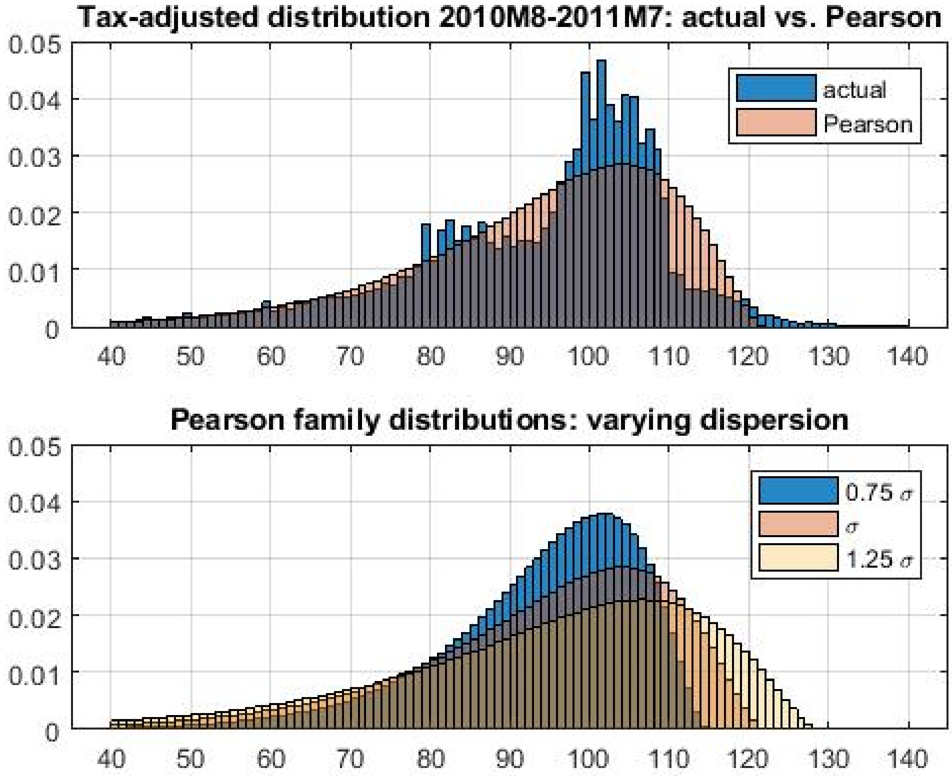

Notes: In the top panel, blue bars represent the Dutch loan-to-value distribution for 2010M8-2011M7 after correcting for the 2011 change in sales taxes through system (1), which we refer to as the "pre-cap distribution". Red bars indicate a Pearson family distribution truncated to (0, 200], which is the same support as the empirical distribution. After truncation, this distribution has essentially the same first four moments as the empirical pre-cap distribution. In the bottom panel, red bars repeat the Pearson family proxy of the pre-cap distribution. Blue bars indicate the "low standard deviation case": a truncated Pearson family distribution which, before truncation and in population, has a standard deviation that is 0.75 that of the proxy of the pre-cap distribution but has the same mean, kurtosis and skewness. Yellow bars indicate the "high standard deviation case", with the same description except for the fact that the population standard deviation before truncation is 1.25 times that of the proxy.

Figure 7: Impulse Responses for Different Degrees of Heterogeneity

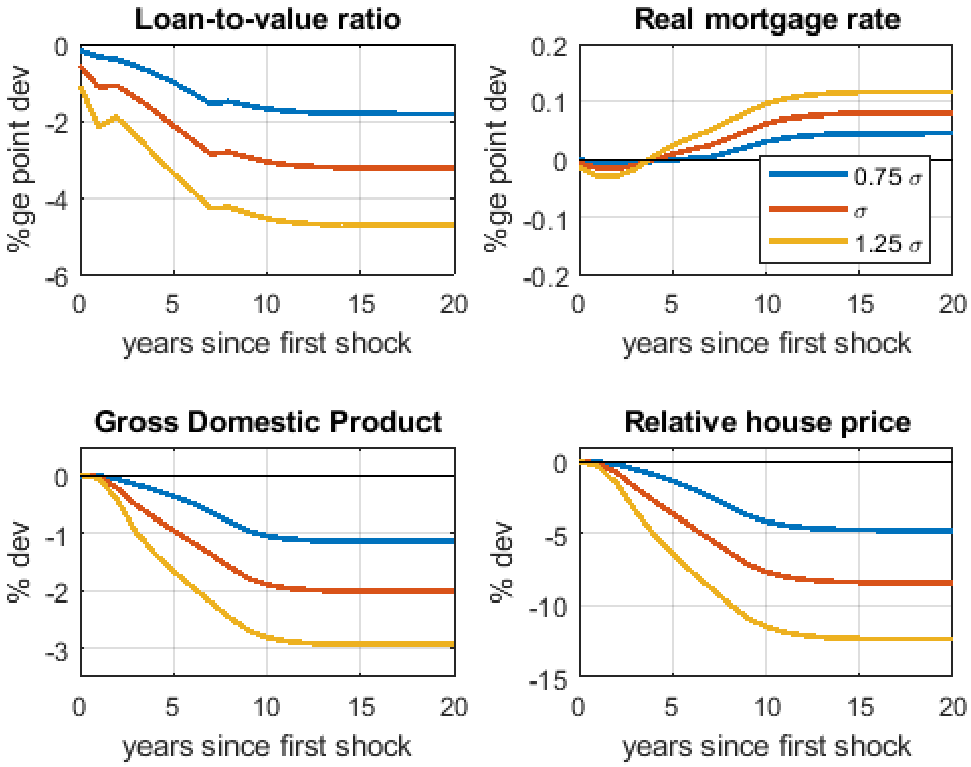

Notes: This figures illustrates that mean-preserving changes in heterogeneity can substantially affect the size of the macro effects of changes in LTV limits. We use the ex ante approach to compute responses to the same changes in the LTV limit and using the same estimated VECM as in Figure 5, but now use different underlying LTV distributions. In the present figure, blue, red and yellow lines pertain to the low, medium and high standard deviation case, respectively, as plotted in Figure 6. In the high standard deviation case, the output and house price responses twenty years after the initial shock are 2.58 times those in the low standard deviation case. Other notes are as for Figure 5.

In the high standard deviation case, house prices are 12.34 percent lower twenty years after the first shock, while GDP is 2.93 percent lower, both relative to the case without a limit. In the low standard deviation case, these responses are 4.79 percent and 1.14 percent, respectively. Both the house price and GDP response are 2.58 times larger in the high standard deviation case than in the low standard deviation case.

We therefore find that if one were to adopt the Dutch LTV policy actions for an economy where the standard deviation of the LTV distribution is 25 percent larger than in the Netherlands, keeping other features of the distribution unchanged, the macroeconomic effects of LTV policy would have been more than twice as large as those in an economy where the standard deviation is 25 percent lower than that in the Netherlands. This illustrates that changes in heterogeneity can have pronounced effects on the size of the macroeconomic effects of changes in LTV limits. To understand why this difference arises, note that for the range of LTV limits which we consider, more dispersion implies that a higher fraction of borrowers is constrained by the limit. As in Section 2, when a higher share of borrowers is constrained, this implies that a higher fraction of individual LTVs move along with changes in the LTV limit, such that the macroeconomic effects are larger.

The results from this section suggest that to measure the effects of a change in the LTV limit, knowing the first four moments of the distribution is sufficient. This is because the responses are similar whether we use the empirical distribution or an approximation for it with the same first four moments. However, knowing the mean of the LTV distribution is not sufficient: the responses are clearly sensitive to the variance of the distribution. In particular, it is important to have information regarding the fraction of households that is constrained by the limit. This point is similar in spirit to that of Krueger, Mitman and Perri (2016), who argue that matching the overall wealth distribution, and in particular matching the share of wealth held by wealth-poor households, is crucial for determining the effect of aggregate shocks on consumption.[31]

[31] Within heterogeneous agent models, it is typically challenging to generate a realistic degree of heterogeneity in the wealth distribution. A measure of such a model's success in matching the empirical wealth distribution is the extent to which it matches the wealth shares of the lowest and highest quintiles of the distribution. In addition to Krueger, Mitman, and Perri (2016), see Carroll, Slacalek, and Tokuoka (2015) and Colciago and Mechelli (2019).

6 Non-Linearity and Asymmetry

In the present section, we investigate non-linearity and asymmetry for a wide range of scenarios for changes in LTV limits. We do this by estimating the effects of hypothetical upward and downward moves in the LTV limit from a range of starting points of the LTV limit. In addition, we investigate the effect of heterogeneity on asymmetry.

The left part of Table 5 reports results for when we apply the ex ante approach to the tax-adjusted Dutch pre-cap distribution of 2010M8-2011M7, again using the VECM from Section 4 to compute impulse responses. In this context, we exclusively consider scenarios where an LTV limit is already in place. The leftmost column indicates the level of the initial limit, ranging from 140 to 50 percent. Each entry in the column "10% up" shows the long-run response, in percent, of the level of relative house prices to a one-off increase in the LTV limit by 10 percentage points from the initial level indicated in that row. The corresponding entry in the column "10% down" shows the effects of a 10 percentage point decrease from that same level. The column "abs ratio" is the ratio of the absolute value of the long-run effect of the downward change to that of the upward change from the same initial cap. This ratio indicates asymmetry from a given starting point.

The responses for real GDP (not shown here) have the same asymmetry ratios, and therefore the same implications for non-linearity and asymmetry as those we discuss for house prices.

The left part of the table yields the following findings.

First, the effects are non-linear. In both the "10% up" and "10% down" columns, the effects grow larger as the starting point for the cap decreases. This non-linearity is reminiscent of that in Section 2.2. It similarly reflects the fact that when the LTV cap is lower relative to the underlying LTV distribution, it is binding for a larger fraction of households, such that changes in the cap have larger macroeconomic effects. At some points, the non-linearity is pronounced. For instance, a tightening in the cap from 110 to 100 percent implies that in the long run, relative house prices are 6.12 percent lower than they would have been without the tightening, while a decrease in the cap from 100 to 90 implies a long-run response of relative house prices by 14.27 percent. This pronounced degree of non-linearity relates to the fact that there is a substantial mass of observations in the range spanned by the two scenarios, i.e. 90-110, such that changes in limits just above 90 percent move substantially more observations than those just below 110 percent.

Second, the effects are asymmetric. From any given starting point, the effect of a decline in the cap is at least as large as that of an increase. For instance, with a starting point of 100 percent, the effect of a ten percentage point decrease in the limit is over twice that of a ten percentage point increase from that same starting point, as indicated by the asymmetry ratio of 2.33. However, asymmetry is minor at levels of the LTV limit that are low relative to the underlying distribution. For instance, a tightening in the LTV limit from 80 to 70 percent implies a response of house prices by 21.58 percent, which in absolute values is only 17 percent larger than the response of an increase from 80 to 90, as indicated by the asymmetry ratio of 1.17. The low degree of asymmetry around an LTV of 80 results from the fact that in that case, the bulk of the mass of the distribution is to the right of the endpoint of the upward move, such that both increases and decreases move the bulk of observations, and therefore have relatively large and quite similar effects.

Symmetry has been the norm in general equilibrium models of the interaction between the housing market and the macroeconomy, including the seminal Iacoviello (2005) model. Our results indicate that in our setting, symmetry fits the data for initial levels of the LTV limit of about 80 and lower, but not so for limits at higher levels. At higher levels, our results favor models with occasionally binding constraints such as Justiniano, Primiceri, and Tambalotti (2015) and Guerrieri and Iacoviello (2017).

Table 5: Asymmetry in the Long-Run Effects of Changes in LTV Limits on House Prices

Initial	Dutch distribution			Initial	Homogeneous LTV		
cap	10% up	10% down	abs ratio	cap	10% up	10% down	abs ratio
140	0.04	-0.07	1.93	140	0.00	0.00	NA
120	0.21	-1.11	5.40	120	0.00	0.00	NA
110	1.11	-6.12	5.50	110	0.00	0.00	NA
105	2.66	-10.72	4.03	105	0.00	0.00	NA
100	6.12	-14.27	2.33	100	0.00	-9.19	$+\infty$
95	10.72	-16.46	1.54	95	0.00	-21.53	$+\infty$
90	14.27	-18.39	1.29	90	9.19	-24.68	2.69
80	18.39	-21.58	1.17	80	24.68	-24.68	1.00
50	23.82	-24.25	1.02	50	24.68	-24.68	1.00

Note: This table documents non-linearity and asymmetry by tabulating the long-run responses of house prices to one-off changes in the LTV limit. The left part of the table, labeled "Dutch distribution", uses the tax-corrected Dutch LTV distribution for 2010M8-2011M7, while the right part, "homogeneous LTV", assumes that all LTVs are at the cross-sectional mean of the Dutch distribution. For both cases, we show the long-run effects on the level of relative house prices of a ten percentage point increase in the LTV limit (column "10% up") and that of a ten percentage point decrease ("10% down"), both from the starting point indicated in the column "initial cap". The column "abs ratio" reports the ratio of the absolute value of the downward effect to that of the upward effect. With the Dutch distribution, we find non-linearity in the sense that changes in the limit from lower starting points have larger effects. We find that asymmetry is important for changes in the limit from starting points of about 90 and up, but not so for starting points of about 80 and below. Comparing the "Dutch distribution" and "homogeneous LTV" cases, we conclude that assuming homogeneity implies that the size of the house price response is substantially underestimated for LTV limits close to but above the average LTV.

Our findings on asymmetry are reminiscent of earlier empirical papers, cited in the introduction, that find that the effect of declines in the LTV cap exceed those of increases. A crucial difference is that in our setting, asymmetry is conditional on the starting point. In Table 5, an increase in the cap that occurs from a low initial level typically has stronger effects than a decrease which occurs from a high initial level. In general, the size of the effects therefore depends on the relative location of the LTV limit and the distribution.

Third, we illustrate the effect of heterogeneity on asymmetry by comparing the above results, based on heterogeneous LTVs, with the case where we assume away any heterogeneity. For the latter, the "homogeneous LTV" case, the right part of Table 5 tabulates long-run house price responses and asymmetry ratios for the same scenarios as those in the left part. In the homogeneous LTV case, we assume that all borrowers have the same LTV of 93.72 percent. That is to say, all borrowers are at the empirical mean of the Dutch pre-cap distribution.

In this case, there is infinite asymmetry for starting points for the cap at 100 and 95, but there is no asymmetry for most initial caps. With homogeneity, non-linearity is pronounced in a small range of the distribution. For instance, a decrease in the limit from 105 to 95 does not affect house prices at all, while a decrease from 90 to 80 implies a house price response of 24.68 percent.

The fact that asymmetry and non-linearity are not present for most initial limits follows from the fact that as long as the limit stays above the homogeneous LTV during any scenario, it is not binding for anyone and therefore has an effect of zero, while as long as the limit stays below that LTV, it is binding for everyone and therefore has the full effect.

The comparison of the homogeneous case with the empirical, heterogeneous case therefore suggests that heterogeneity tends to spread out asymmetry and non-linearity over a wider range.

These differences in the shape of asymmetry are such that if one were to assume homogeneous LTVs, one would understate the effects of changes in the LTV limit above the mean of the underlying distribution. With the Dutch distribution, an increase in the LTV limit from 95 to 105 implies that house prices are 10.72 percent higher than they would have been if the limit had stayed at 95, while the homogeneous LTV setting implies no effect.

In sum, we find that if one abstracts from borrower heterogeneity, one overstates asymmetry around the cross-sectional average for the LTV and therefore understates the size of the effects of LTV limits when these lie above the average LTV. This is relevant for homogeneous LTV models with occasionally binding credit constraints such as Justiniano, Primiceri, and Tambalotti (2015) and Guerrieri and Iacoviello (2017).[32]

[32] Our analysis does not replicate either of these models. For instance, unlike what is the case for our analysis, the

7 Conclusion

In this paper, we implemented a new approach to estimating the macroeconomic effects of changes in loan-to-value limits. While regression-based analysis in earlier literature assumes that these effects are the same across countries and across episodes, we estimate responses for a single episode in a single economy. Our approach accounts for the position and shape of the underlying distribution of loan-to-value ratios, i.e. the distribution that would obtain if there had been no cap. In this sense, it accounts for the repercussions on this distribution of factors including mortgage credit cycles, other macroprudential policy measures and changes in household income expectations.

Our main results are as follows. First, our estimates suggest that loan-to-value policy actions in the Netherlands in 2011-2018 have a non-trivial effect on relative house prices and real GDP. When using a comparable policy scenario, our results are similar to those in our contribution to De Nederlandsche Bank (2015a). We do not assess whether loan-to-value policy has been welfare enhancing and we do not evaluate what is the optimal level of the LTV limit. Unlike the study we just mentioned, we do not discuss the effects of LTV limits on financial stability.

Second, we find that changes in the LTV cap have non-linear effects in the sense that the effects become larger as the cap tightens. This occurs because the fraction of households for which the cap binds increases as the cap tightens.

Third, we document that a moderate, mean-preserving change in the degree of heterogeneity can substantially alter the effects of changes in LTV caps. This highlights the importance of taking

two models explicitly account for the effects of expected future credit constraints. In Guerrieri and Iacoviello (2017), asymmetry around the steady state is not infinite. The reason is that even when the current collateral constraint is not binding for the borrower, future constraints are, such that a further easing in the current constraint tends to come with a relaxation in the extent to which these future constraints are expected to be binding, and therefore still has some positive effect on borrower spending.

heterogeneity in credit markets into account when gauging the effect of changes in loan-to-value limits. This is important because existing empirical literature on measuring the effects of changes in loan-to-value limits typically focuses on the first moment of the LTV distribution.

Fourth, from a given initial level for the LTV cap and for a given underlying LTV distribution, a tightening tends to affect more borrowers than an easing, and therefore tends to have stronger effects.

Fifth, we find that with the LTV distribution that prevailed in the Netherlands, symmetry is a pretty good approximation for changes in LTV limits around 80 percent or lower since in that area the limit is binding for the bulk of households. Asymmetry prevails for higher starting points for LTV limits.

Sixth, when we impose homogeneity, there is extreme asymmetry for a narrow range of initial LTV limits, whereas asymmetry is more spread out in a heterogeneous world. This implies that when the cap exceeds the average LTV, the effects of changing the limit are larger if one accounts for heterogeneity than if one assumes that all borrowers have the same LTV.

Overall, our results imply that a change in an LTV limit has stronger macro effects when a higher fraction of households is constrained by the limit. We see the following implications of this feature.

First, by documenting non-linearity, we illustrate that changes in the tightness of credit constraints can alter the fraction of credit-constrained households to such an extent that these changes in credit constraints have a stronger effect on the economy. While we are not aware of any other papers documenting this type of non-linearity, it would be more naturally captured in a heterogeneous agents model than in a homogeneous borrower model.

Second, consider a financial authority that wishes to set an LTV limit to a particular level with the intention of keeping it constant, e.g. so as to enhance financial stability. Our results suggest that the immediate macro effects of imposing or tightening a limit are smaller when there are relatively few high LTVs. Such a situation could arise, for instance, if other aspects of financial regulation make it more costly for banks to hold high-LTV mortgages, or if households desire to borrow at lower loan-to-value ratios because they expect lower house price inflation.

In terms of an avenue for further research, note that in the implementation of our approach in this paper, we use a recent loan-to-value distribution as a proxy for what the current distribution would be if there were no loan-to-value cap or if the cap had been less tight. This type of proxy is not feasible in all settings. In particular, when one wants to study a loosening in the LTV limit but there is no recent pre-cap distribution available as a measure for the underlying distribution, one would need to construct an estimate for the underlying LTV distribution in another manner.

Appendix

This appendix provides an English translation of survey questions we used to construct our loan-to-value and loan-to-income series from Section 3. Square brackets indicate words that enter conditional on answers to previous questions. Accolades indicate comments to readers of this appendix that did not enter the survey.

Introduction

This questionnaire concerns home purchases. We request that you fill in the questions, whether you ever bought a home or not.

Did you ever buy a home? Then you can expect to find questions about your **first** home purchase. You may still be living there; it is also possible that you have moved (several times) in the mean time. Do you still have documents with information about the purchase of this first home? In that case we request that you take the documents from which you can tell the date of purchase, the purchase price and the mortgage amount. If you do not have these documents, we kindly request that you give your best estimate for the relevant questions.

v0 Thank you for filling in this questionnaire. Have you ever bought a home (a house or apartment)?

1 Yes 2 No

if v0=1

v02 As indicated before, this questionnaire is about your **first** home purchase. You may still be living there; it is also possible that you have moved (several times) in the mean time.

Did you buy your first home by yourself or together with someone else? We are looking for the buyer(s) in the purchase contract.

1 By myself 2 Together with my current partner 3 Together with a former partner 4 Other

if v02=2 or v02=3

v03 You bought your first home together with your [*if v02=3*: former] partner. As for your [*if v02=3*: former] partner, was it also the first time for him or her to buy a home?

1 Yes, it was the first time to buy a home for both of us

2 No, my [*if v02=3*: former] partner had bought a house before

99 I don't know

if v0=1

v04_t1 - v04_t2

When did you buy your first home? We mean the moment when you signed the purchase agreement. Please indicate the year and, if you still remember, the month of signature.

v04_t1 Year 1910..2016

v04_t2 Month 1 January 2 February 3 March 4 April 5 May 6 June 7 July 8 August 9 September 10 October 11 November 12 December 13 I don't know

if v0=1

v05_t1 - v05_t3 What was the purchase price at which you bought your first home? By purchase price of a home, we mean the amount that you agreed upon with the seller. Please do not include any 'kosten koper'{ additional costs paid by the buyer}, such as transaction taxes or notary fees, in the purchase price.

If you do not remember the purchase price, we request that you give your best estimate.

Please fill in a whole number without spaces, commas or points.

v05_t1 Amount

v05_t2 1 euro 2 guilder

v05_t3 I cannot give an estimate

0 No 1 Yes

if v0=1

v06 Did you take out a mortgage loan in connection with the purchase of your first home?

1 Yes 2 No

if v06=1

v07 How did the total mortgage amount compare to the purchase price? If you took out more than one mortgage loan in connection with the purchase of your first home, we request that you sum the amounts of these loans. We ask you once more to not include any 'kosten koper' in the purchase price.

The total amount of the mortgage loan was

1 more than 30% above the purchase price
2 about 30% above the purchase price
3 about 20% above the purchase price
4 about 10% above the purchase price
5 about 5% above the purchase price
6 about equal to the purchase price
7 about 5% below the purchase price
8 about 10% below the purchase price
9 about 20% below the purchase price
10 about 30% below the purchase price
11 about 40% below the purchase price
12 about 50% below the purchase price
13 more than 50% below the purchase price

if v06=1 and v05_t3 $\neq$ *1*

v08_t1 - v08_t3 You bought your first home for a purchase price of [v05_t1] [v05_t2], without including 'kosten koper'. Do you still remember the total mortgage amount that you borrowed when you purchased this home? If so, please fill in this total mortgage amount. We ask you once more to sum the borrowed amounts in case you took out multiple mortgage loans at that time.

If you do not remember the total mortgage amount anymore, we request that you give your best estimate here too.

Please fill in a whole number without spaces, commas or points.

v08_t1 Amount

v08_t2 1 euro 2 guilder

v08_t3 I cannot provide an estimate

0 No 1 Yes

if (v02=2 or v02=3) and v06=1
v12 You indicated that you bought your first home with your [*if v02=3*: former] partner.
Which income(s) did the bank or other credit provider take into account when assessing your application for a mortgage loan?
1 Only my income 2 Only the income of my [*if v02=3*: former] partner 3 Both my income and that of my [*if v02=3*: former] partner

if (v02=1 and v06=1) or v12=1
v13a
Approximately how did the total mortgage amount compare to your total annual gross income in the year you purchased the home? The total mortgage amount was ... times my total annual gross income.
1 0.5
2 1
3 1.5
4 2
5 2.5
6 3
7 3.5
8 4
9 4.5
10 5
11 5.5
12 6
13 6.5
14 7
15 7.5
16 8
17 More than 8

if v12=3
v13b Approximately how did the total mortgage amount compare to your total annual gross family income in the year you purchased the home? The total mortgage amount was ... times our total annual gross family income in the year we purchased the home.
1 0.5
2 1
3 1.5
4 2
5 2.5
6 3
7 3.5
8 4
9 4.5
10 5
11 5.5
12 6
13 6.5
14 7

15 7.5
16 8
17 More than 8

if v12=2
v13c Do you know how the total mortgage approximately compared to the total gross annual income of your [*if v02=3*: former] partner in the year you purchased your first home? The total mortgage amount was ... times the total annual gross income of my partner.
1 0.5
2 1
3 1.5
4 2
5 2.5
6 3
7 3.5
8 4
9 4.5
10 5
11 5.5
12 6
13 6.5
14 7
15 7.5
16 8
17 More than 8
99 I don't know

if ((v02=1 and v06=1) or (v12=1) or (v12=3))
v14_t1 - v14_t3
Are you still in a position to give a good estimate of your annual gross [*if ((v02=1 and v06=1) or (v12=1))*: income/ *if (v12=3))*: family income] in the year your purchased your first home? If so, what was your total gross [*if ((v02=1 and v06=1) or (v12=1))*: income/ *if (v12=3))*: family income] in the year of purchase?
Please fill in a whole number without spaces, commas or periods.
v14_t1 Amount
v14_t2 1 euro 2 guilder
v14_t3 I cannot give an estimate
0 No 1 Yes

References

[1] Akinci, Ozge, and Jane Olmstead-Rumsey. (2018) "How Effective Are Macroprudential Policies? An Empirical Investigation." *Journal of Financial Intermediation*, 33, 33-57.

[2] Almeida, Heitor, Murillo Campello, and Crocker Liu. (2006) "The Financial Accelerator: Evidence from International Housing Markets." *Review of Finance*, 10, 321-352.

[3] Brzoza-Brzezina, Michal, Paolo Gelain, and Marcin Kolasa. (2014) "Monetary and Macroprudential Policy with Multi-Period Loans." Norges Bank Working Paper 16/2014.

[4] Bussel, Arjan, P. Kerkhoffs and Ronald Mahieu. (1996) "Een Nieuwe Index voor de Huizenmarkt." *Economisch Statistische Berichten*, 81, 897-899.

[5] Calza, Alessandro, Tommaso Monacelli, and Livio Stracca. (2013) "Housing Finance and Monetary Policy." *Journal of the European Economic Association*, 11, 101-122.

[6] Carroll, Christopher D., Jiri Slacalek, and Kiichi Tokuoka. (2015) "Buffer-stock Saving in a Krusell-Smith World", *Economics Letters*, 132, 97-100.

[7] Catte, Pietro, Nathalie Girouard, Robert Price, and Christophe André. (2004) "Housing Markets, Wealth and the Business Cycle.", Organisation for Economic Co-operation and Development, Economics Department Working Paper No. 394.

[8] Cerutti, Eugenio, Stijn Claessens, and Luc Laeven. (2017) "The Use and Effectiveness of Macroprudential Policies: New Evidence." *Journal of Financial Stability*, 28, 203-224.

[9] Colciago, Andrea, and Rajssa Mechelli. (2019) "Competition and Inequality: Aiyagari meets Bertrand and Cournot". University of Milano - Bicocca, DEMS Working Paper No. 398.

[10] CPB Netherlands Bureau for Economic Policy Analysis. (2015) "Economische Effecten van Verdere Verlaging LTV-Limiet". Note to the Dutch Financial Stability Committee.

[11] CPB Netherlands Bureau for Economic Policy Analysis. (2017) "Actualisatie Economische Effecten van een Verdere Verlaging van de LTV-Limiet". CPB Background Document, 6 juni 2017.

[12] Crowe, Christopher, Giovanni Dell'Ariccia, Deniz Igan, and Pau Rabanal. (2013) "How to Deal with Real Estate Booms: Lessons from Country Experiences." *Journal of Financial Stability*, 9, 300-319.

[13] Debortoli, Davide, and Jordi Galí. (2017) "Monetary Policy with Heterogeneous Agents: Insights from TANK Models." Mimeo, Universitat Pompeu Fabra.

[14] De Nederlandsche Bank. (2015a) "Effects of Further Reductions in the LTV Limit." De Nederlandsche Bank Occasional Studies 13-2.

[15] De Nederlandsche Bank. (2015b) "Dutch Mortgages in the DNB Loan Level Data." De Nederlandsche Bank Occasional Studies 13-4.

[16] Duca, John V., John Muellbauer, and Anthony Murphy. (2011) "House Prices and Credit Constraints: Making Sense of the U.S. Experience." *Economic Journal*, 121, 533-551.

[17] Duca, John V., John Muellbauer, and Anthony Murphy. (2012) "Shifting Credit Standards and the Boom and Bust in U.S. House Prices: Time Series Evidence from the Past Three Decades." Mimeo, Federal Reserve Bank of Dallas.

[18] Elliott, Douglas J., Greg Feldberg, and Andreas Lehnert. (2013) "The History of Cyclical Macroprudential Policy in the United States." U.S. Department of the Treasure, Office of Financial Research Working Paper No. 8.

[19] European Commision. (2014) "Tax Reforms in EU Member States. Tax Policy Challenges for Economic Growth and Fiscal Sustainability." European Economy 6|2014.

[20] Galati, Gabriele, and Richhild Moessner. (2017) "What Do We Know About the Effects of Macroprudential Policy?" *Economica*, 1-36.

[21] Gerlach, Stefan, and Wensheng Peng. (2005) "Bank Lending and Property Prices in Hong Kong." *Journal of Banking and Finance*, 29, 461-481.

[22] Glaeser, Edward L., Joshua D. Gottlieb, and Joseph Gyourko. (2013) "Can Cheap Credit Explain the Housing Boom?" In *Housing and the Financial Crisis*, edited by Edward L. Glaeser and Todd Sinai, 301-359. Cambridge, MA: National Bureau of Economic Research.

[23] Guerrieri, Luca, and Matteo Iacoviello. (2017) "Collateral Constraints and Macroeconomic Asymmetries." *Journal of Monetary Economics*, 90, 28-49.

[24] Guerrieri, Veronica, and Guido Lorenzoni. (2017) "Credit Crises, Precautionary Savings, and the Liquidity Trap." *Quarterly Journal of Economics*, 132, 1427-1467.

[25] Iacoviello, Matteo. (2005) "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle." *American Economic Review*, 95, 739-764.

[26] Igan, Deniz, and Heedon Kang. (2011) "Do Loan-to-Value and Debt-to-Income Limits Work? Evidence from Korea." IMF Working Paper No. 11/297.

[27] Johansen, Soren. (1988) "Statistical Analysis of Cointegration Vectors." *Journal of Economic Dynamics and Control*, 12, 231-254.

[28] Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti. (2015) "Household Leveraging and Deleveraging." *Review of Economic Dynamics*, 18, 3-20

[29] Kaplan, Greg, and Giovanni L. Violante. (2018) "Microeconomic Heterogeneity and Macroeconomic Shocks." *Journal of Economic Perspectives*, 32, 167-194.

[30] Krueger, Dirk, Kurt Mitman, and Fabrizio Perri. (2016) "Macroeconomics and Household Heterogeneity." In *Handbook of Macroeconomics,* edited by John B. Taylor and Harald Uhlig, pp. 843-921. Amsterdam, the Netherlands: Elsevier.

[31] Kuttner, Kenneth N. and Ilhyock Shim. (2016) "Can Non-Interest Rate Policies Stabilize Housing Markets? Evidence from a Panel of 57 Economies." *Journal of Financial Stability*, 26, 31-44.

[32] Lamont, Owen, and Jeremy C. Stein. (1999) "Leverage and House-Price Dynamics in U.S. Cities", *RAND Journal of Economics*, 30, 498-514.

[33] Ling, David C., Joseph T.L. Ooi, and Thao T.T. Le. (2015) "Explaining House Price Dynamics: Isolating the Role of Nonfundamentals." *Journal of Money, Credit and Banking*, 47, 87-125.

[34] McDonald, Chris. (2015) "When is Macroprudential Policy Effective?" In *Macroprudential Policy and Practice*, edited by Paul Mizen, Margarita Rubio and Philip Turner, 248-274. Cambridge University Press.

[35] Shim, Ilhyock, Bilyana Bogdanova, Jimmy Shek, and Agne Subelyte. (2013) "Database for Policy Actions on Housing Markets." BIS Quarterly Review, September 2013, 83-95.

[36] Timmermans, Maud. (2012) "The Effects of Tightening the Mortgage Loan-to-Value Ratio in the Netherlands", Master's Thesis, University of Amsterdam.

[37] Wong, Eric, Tom Fong, Ka-fai Li, and Henry Choi. (2011) "Loan-to-Value Ratio as a Macro-Prudential Tool - Hong Kong's Experience and Cross-Country Evidence." Hong Kong Monetary Authority Working Paper 01/2011.

[38] Zdzienicka, Aleksandra, Sally Chen, Federico Diaz Kalan, Stefan Laseen, and Katsiaryna Svirydzenka. (2015) "Effects of Monetary and Macroprudential Policies on Financial Conditions: Evidence from the United States." IMF Working Paper No. 15/288.

Previous DNB Working Papers in 2019

No. 622 **David-Jan Jansen,** Did Spillovers From Europe Indeed Contribute to the 2010 U.S. Flash Crash?
No. 623 **Wilko Bolt, Kostas Mavromatis and Sweder van Wijnbergen,** The Global Macroeconomics of a trade war: the EAGLE model on the US-China trade conflict
No. 624 **Ronald Heijmans and Chen Zhou,** Outlier detection in TARGET2 risk indicators
No. 625 **Robert Vermeulen, Edo Schets, Melanie Lohuis, Barbara Kölbl, David-Jan Jansen and Willem Heeringa,** The Heat is on: A framework measuring financial stress under disruptive energy transition scenarios
No. 626 **Anna Samarina and Anh D.M. Nguyen,** Does monetary policy affect income inequality in the euro area?
No. 627 **Stephanie Titzck and Jan Willem van den End,** The impact of size, composition and duration of the central bank balance sheet on inflation expectations and market prices
No. 628 **Andrea Colciago, Volker Lindenthal and Antonella Trigari,** Who Creates and Destroys Jobs over the Business Cycle?
No. 629 **Stan Olijslagers, Annelie Petersen, Nander de Vette and Sweder van Wijnbergen,** What option prices tell us about the ECB's unconventional monetary policies
No. 630 **Ilja Boelaars and Dirk Broeders,** Fair pensions
No. 631 **Joost Bats and Tom Hudepohl,** Impact of targeted credit easing by the ECB: bank-level evidence
No. 632 **Mehdi El Herradi and Aurélien Leroy,** Monetary policy and the top one percent: Evidence from a century of modern economic history
No. 633 **Arina Wischnewsky, David-Jan Jansen and Matthias Neuenkirch,** Financial Stability and the Fed: Evidence from Congressional Hearings
No. 634 **Bram Gootjes, Jakob de Haan and Richard Jong-A-Pin,** Do fiscal rules constrain political budget cycles?

DeNederlandscheBank

EUROSYSTEEM

De Nederlandsche Bank N.V.
Postbus 98, 1000 AB Amsterdam
020 524 91 11
dnb.nl