Do Macroprudential Policies Stabilize Consumption in a Downturn?

Evidence from administrative data on leverage, spending, and monetary policy transmission*

Knut Are Aastveit, Ragnar Enger Juelsrud, Karolis Liaudinskas,
José-Luis Peydró and Ella Getz Wold

March 2026

PRELIMINARY AND INCOMPLETE - PLEASE DO NOT DISTRIBUTE

Abstract

Does macroprudential regulation stabilize household consumption during economic downturns, and how does it interact with monetary policy? We answer these questions by exploiting a tightening of the loan-to-value (LTV) cap in Norway from 90% to 85% in 2012, a subsequent regional recession triggered by the 2014 collapse in oil prices, and identified monetary policy shocks, combined with household-level administrative data on income, wealth, housing transactions, and consumption. Comparing prospective first-time home buyers immediately before and after the policy change, we document three findings. First, households exposed to the tighter LTV cap enter the recession with substantially lower leverage, driven primarily by a 28% reduction in the probability of homeownership rather than by changes in leverage among buyers. Second, tighter macroprudential policy during the boom leads to significantly smaller consumption declines during the recession. While consumption in the affected region falls by about 8% relative to the rest of Norway, the decline is only around 3% for households subject to the stricter cap, corresponding to a 60% attenuation. Using a principal stratification framework, we show that this stabilization is almost entirely driven by the extensive margin. Third, tighter macroprudential policy weakens the transmission of monetary policy to consumption during the downturn. Our results reveal a resilience–effectiveness tradeoff: macroprudential policy makes monetary policy less powerful, but also less necessary. More broadly, we provide causal evidence that ex-ante macroprudential tightening stabilizes household demand in an actual recession.

JEL-codes: E21, E58, G01, G21, G28, G51.
Keywords: Macroprudential policy, Monetary Policy, Leverage, Consumption, Recession.

---

*This paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. Knut Are Aastveit is at Norges Bank and BI Norwegian Business School; Ragnar Enger Juelsrud is at University of Oslo and CEPR; Karolis Liaudinskas is at Norges Bank; Jose-Luis Peydro is at LUISS University, EIEF, and CEPR; and Ella Getz Wold is at the BI Norwegian Business School.

1 Introduction

Do macroprudential policies stabilize household consumption during economic downturns, and do they affect the transmission of monetary policy to consumption? Borrower-based macroprudential regulations, such as loan-to-value (LTV) caps, are designed to curb household leverage during expansions, thereby limiting consumption declines when recessions occur (Bernanke, 2011). Although a growing literature studies the immediate effects of LTV regulation, there is little empirical evidence on whether these policies achieve their ultimate goal of stabilizing consumption in recessions, or on how they interact with monetary policy. A key insight from the theoretical literature is that macroprudential policy, by reshaping household balance sheets, can also influence the transmission of monetary policy during downturns (e.g., Farhi and Werning 2016; Van der Ghote 2021; Martín et al. 2021). We provide evidence on both questions by studying how a tightening of LTV caps affects household consumption during a subsequent regional recession, and how it shapes the transmission of monetary easing. We use a difference-in-differences design to identify the causal effects of macroprudential policy, and local projections to study how these balance sheet differences shape the dynamic response to monetary policy shocks. More broadly, our paper shows how macroprudential policy creates ex-ante differences in household balance sheets that shape the economy's ex-post response to aggregate shocks. Answering these questions requires a setting with exogenous macroprudential variation, a realized economic downturn, and identified monetary policy shocks.

Norway provides an excellent setting for studying these questions. We exploit three sources of variation. First, a tightening of the LTV cap from 90% to 85%, announced in December 2011 and effective from January 2012, created a sharp discontinuity in borrowing constraints facing prospective home buyers. Second, the collapse in international oil prices in 2014 triggered a regional recession concentrated in Rogaland County, where petroleum activities account for 25–30% of employment compared to less than 5% nationally. Household consumption in Rogaland fell approximately 8% relative to the rest of Norway. This regional variation allows us to compare consumption responses to the same aggregate shock across households with different policy-induced leverage. Third, we observe identified monetary policy shocks during the recession period, enabling us to study how macroprudential regulation shapes the transmission of monetary easing. The analysis draws on an exceptionally rich data environment: household-level administrative data linking tax records (income and wealth), the Land Registry (housing transactions), and debit card transactions covering approximately 80% of all card payments in Norway.

We compare two cohorts of prospective first-time home buyers: households that were on the margin of purchasing in 2011, when the LTV cap was 90%, and analogous households in 2012, when the cap had tightened to 85%. A household qualifies as a prospective buyer if it does not yet own a home but has accumulated sufficient liquid wealth to meet the prevailing down payment requirement. Our identification strategy is a triple difference-in-differences design, comparing consumption growth for treated versus control cohorts, in the recession region (Rogaland) versus the

rest of Norway, before versus during the downturn. We control for household fixed effects, time or municipality-by-time fixed effects, and initial balance sheet characteristics interacted with the recession period. Importantly, the two cohorts are balanced on observables prior to the policy change.

Results. We present three sets of main findings. First, the tightening of the LTV cap reduces leverage primarily through the extensive margin. The policy change reduces the probability of home purchase by 4.0 percentage points, from 14.5% to 10.5%, representing a 28% decline relative to the control group mean. Among the entire sample of prospective buyers, total debt falls by 21% for the treatment group. This decline operates almost entirely through fewer households purchasing homes rather than through lower leverage among those who do purchase.

Among households that do purchase, we observe clear bunching in the LTV distribution: the mass of buyers with LTV ratios between 85% and 90% drops sharply, shifting to ratios just below 85%. However, the median LTV among buyers changes little and the aggregate debt reduction conditional on buying is small. The extensive margin dominates: deterring 28% of prospective buyers from purchasing generates a far larger aggregate effect on total debt than the modest leverage adjustments among those who still buy.

These balance sheet differences persist into the recession period. By 2014, treated households have 17% lower housing wealth, 8% lower debt, and are 11% less likely to be homeowners. Debt-to-income ratios are 5.3% lower and debt-to-wealth ratios are 3.3% lower for the treatment group. The policy-induced differences in leverage remain large and economically meaningful at the onset of the downturn.

Second, tighter macroprudential policy during the boom leads to substantially smaller consumption declines during the recession. Consumption in Rogaland falls approximately 8% relative to the rest of Norway. For households exposed to the stricter LTV cap, however, the decline is only about 3%. The 5 percentage-point treatment effect is statistically significant and economically large, representing a roughly 60% reduction in the consumption decline. We verify parallel pre-trends: estimating the triple-difference specification year by year with 2014 as the omitted category, the treatment coefficients prior to the recession are small and statistically insignificant.

The result is robust to a wide range of specifications, including household fixed effects, municipality-by-year fixed effects, controls for initial balance sheet positions and sector-specific income shares, and log consumption as the dependent variable. We also show that the treatment and control groups do not differ in their exposure to oil-sector income risk or in the probability of receiving unemployment insurance during the downturn, ruling out differential income shocks as a driver of the consumption results.

We further show that the consumption stabilization concentrates in the categories most affected by the recession. We decompose total consumption into six groups (food, services, recreation, personal goods, utilities and housing, and transportation). Transport and utilities — which include

vehicle purchases, furniture, and household equipment — decline most sharply during the recession, and these are precisely the categories where the treatment effect is largest. The correlation between recession effects and treatment effects across consumption groups is 0.85, providing within-household validation that the stabilization occurs where the downturn hits hardest.

To better understand the mechanism, we decompose the total consumption effect into contributions from the extensive margin – households deterred from purchasing a home, and the intensive margin, namely balance sheet adjustments among those who still buy. Using a principal stratification framework, we show that, under the maintained assumptions, the consumption stabilization effect is driven primarily by the extensive margin, while the intensive margin plays a limited quantitative role. This finding has important policy implications. The consumption smoothing benefits of tighter LTV caps do not arise from improving the resilience of leveraged homeowners, but from preventing marginal households from entering homeownership in the first place. In other words, macroprudential policy stabilizes consumption not by making borrowers safer, but by reducing the number of borrowers.

Third, our results suggest that macroprudential regulation reduces both the need for and the effectiveness of monetary policy. We examine how the LTV tightening affects the transmission of monetary policy to consumption during the recession, using narrative monetary policy shocks from Holm et al. (2021) and high-frequency shocks from Brubakk et al. (2022). An expansionary monetary policy shock increases consumption for control households during the recession, but the effect is substantially attenuated for treated households. The difference is both economically and statistically significant: monetary policy is substantially less powerful for households with lower leverage.

At the same time, treated households require less monetary stimulus. Their baseline consumption decline is 5 percentage points smaller. Control households experience an 8% consumption decline and respond strongly to monetary easing; treated households experience only a 3% decline but respond weakly. The consumption gap that monetary policy must close is therefore smaller for treated households, even though the policy instrument is less effective at closing it.

These results reveal a resilience–effectiveness tradeoff: macroprudential policy dampens household demand fluctuations, but also reduces the sensitivity of consumption to monetary policy. The welfare implications depend on the constraints facing the central bank: near the effective lower bound, the reduced need for stimulus is particularly valuable. Our evidence suggests that macroprudential tightening may reduce the probability of hitting such constraints by dampening the initial consumption decline.

Contribution to the literature We make two main contributions. First, we provide, to our knowledge, the first causal evidence that ex-ante macroprudential tightening stabilizes household consumption during a realized economic downturn. Second, we document that macroprudential regulation weakens the transmission of monetary policy during downturns, pointing to a resilience–

effectiveness tradeoff. We also contribute to the literature on the effects of LTV requirements on household balance sheets.

First, we contribute to the literature on the "credit-driven household demand channel," which argues that severe recessions are often precipitated by a collapse in spending among over-leveraged households, e.g. Mian and Sufi (2011); Mian et al. (2013, 2017); Carroll et al. (2011) on the empirical side and Jones et al. (2022); Kaplan et al. (2018) on the theoretical side. For instance, Mian et al. (2013) demonstrate that US counties with higher pre-crisis leverage experienced consumption declines 2–3 times larger than low-leverage counties during the Great Recession. Mian et al. (2017) extend this to 30 economies over 40 years, finding that a one standard deviation increase in household debt-to-GDP predicts a 2.1 percentage point decline in subsequent three-year GDP growth. Jones et al. (2022) quantify these effects using a heterogeneous agent model, finding that household deleveraging explains 40% of the differential employment changes across US states during the Great Recession. This provides theoretical grounding for expecting that lower initial leverage — from, e.g., LTV caps — would reduce deleveraging pressure during downturns.

Despite extensive evidence that leverage amplifies consumption declines, no prior study provides causal evidence that ex-ante macroprudential policy stabilizes consumption during an actual, realized economic downturn. Existing macroprudential effectiveness research focuses on crisis prevention — reducing the probability of adverse events — mostly using country-level data, rather than on consumption smoothing during crises. Other papers, e.g. Jiménez et al. (2017), study firm-level outcomes but not household consumption. Bergant et al. (2023) examine whether pre-pandemic macroprudential tightening reduced economic stress during COVID-19, finding that countries that tightened more aggressively experienced less financial stress, but their focus is financial stress metrics rather than household consumption. We contribute to this strand of the literature by providing, to our knowledge, the first causal evidence that ex-ante macroprudential intervention successfully mitigates consumption declines during a recession.

Second, we contribute to the literature on the interaction between macroprudential and monetary policy. Several theoretical studies outline the principles behind coordinating these tools. Van der Ghote (2021) shows that when macroprudential policy optimally addresses systemic risk, the Tinbergen principle applies — monetary policy can focus narrowly on price stability. Without effective macroprudential policy, optimal monetary policy must deviate from inflation stabilization to address financial imbalances. Martín et al. (2021) extend this analysis, finding that effective macroprudential policy reduces the burden on monetary policy to address financial stability concerns. International Monetary Fund (2013) establishes that macroprudential policies can attenuate the "side effects" of monetary policy on financial stability.

The empirical literature demonstrates the importance of household balance sheets for the transmission of monetary policy, e.g. Flodén et al. (2021); Cloyne et al. (2020); Holm et al. (2021). As such, it is plausible that macroprudential measures such as LTV caps, which affect household bal-

ance sheets, also affect monetary transmission. At the aggregate level, International Monetary Fund (2024) finds that the transmission of monetary policy is stronger when fixed-rate mortgages are less common, home buyers are more leveraged due to less restrictive LTV limits, and household debt is high. Bank of Canada (2024) documents that tighter macroprudential policies during low-rate periods helped limit risk-taking but may have come at the cost of muting some expansionary effects of loose monetary policy. On the other hand, Alpanda and Zubairy (2019) find that monetary policy effects are less powerful during periods of high household debt — the home equity loan channel weakens when initial debt levels are elevated — while Greenwald (2018) shows LTV caps may make monetary policy more effective at stabilizing inflation while dampening credit transmission. Harding and Klein (2022) find that monetary policy has stronger effects when borrowing constraints bind and household net worth is low.

We contribute to this literature by providing evidence on a resilience–effectiveness tradeoff in the interaction between macroprudential and monetary policy. Our results suggest that stricter macroprudential regulation makes monetary policy less effective but also less necessary during downturns. Because macroprudential policy stabilized consumption by reducing ex-ante leverage, the gap that monetary policy must close is smaller during recessions. This effectively reduces the potency of the central bank's primary tool while simultaneously lowering the tail risk of hitting the effective lower bound during a crisis. Our evidence suggests that the benefits of pre-crisis macroprudential tightening — in terms of reduced consumption volatility and smaller required monetary response — may outweigh the costs of somewhat reduced monetary policy potency.

Finally, we contribute to the literature on the impact of LTV caps on household credit and other balance sheet variables. Initially, much of this literature focused on macro-level analyses (Cerutti et al., 2017; Kuttner and Shim, 2016; International Monetary Fund, 2017). Recently, several micro-level studies have provided evidence on balance sheet adjustments at the household level. van Bekkum et al. (2024) use Dutch administrative data to show that LTV limits induce affected households to experience less financial distress. DeFusco et al. (2020) employ regression discontinuity around Dodd-Frank leverage rules, demonstrating that regulation eliminated 15% of the affected market entirely and reduced leverage for another 20% of borrowers. Acharya et al. (2020) study Ireland's 2015 LTV/LTI limits, finding credit reallocation from low- to high-income borrowers that weakens the credit–house price feedback loop. A closely related paper to ours is Aastveit et al. (2024), which analyzes the impact of LTV restrictions in Norway, primarily at the intensive margin. They find that while LTV caps reduce leverage, they also reduce liquidity due to higher down payment requirements.

We contribute to this literature by identifying the extensive margin as the dominant driver of balance sheet improvements, and — more importantly — by analyzing the effects during a recession. While previous studies have focused primarily on the intensive margin (the debt levels of those who still buy a home), we show that the main effect of the 2012 Norwegian LTV tightening

was a 28% reduction in the probability of purchasing a home among potential first-time buyers. Consequently, the average potential buyer entered the 2014 recession with 21% less debt and higher aggregate liquid wealth. This suggests that macroprudential policy stabilizes demand primarily not through fine-tuning the leverage of homeowners, but by preventing marginal households from purchasing. Moreover, we show that tighter macroprudential policy during the boom leads to smaller consumption declines during the recession.

2 Institutional setting

Our analysis exploits a tightening of Norwegian macroprudential policy in 2012 and a regional economic downturn caused by the 2014 oil price collapse. We describe each in turn.

2.1 Mortgage regulation

In 2011, Norwegian borrowers faced an LTV cap of 90%. In December 2011, it was announced that the cap would be tightened to 85%, effective from January 2012. As a result, prospective home buyers in 2012 faced stricter borrowing constraints than those in 2011.

It is important to note that the LTV caps were not hard caps, in the sense that banks were given room to deviate. The LTV on a given mortgage could exceed the cap if the bank could document that an extraordinary risk assessment had taken place. The room to deviate was later formalized into a flexibility quota, allowing 8–10% of new mortgages each quarter to exceed the cap.

As is common, the LTV cap applied to mortgage debt connected to the specific house purchase and did not restrict total household debt. It was therefore possible for a household's total debt relative to house value to exceed the cap. Households with parental support also had avenues to circumvent the constraint — for instance, if a parent pledged collateral in the form of their own property or agreed to be jointly liable for the mortgage.

All in all, the implementation of the LTV regulation allowed a non-trivial share of mortgages to have LTV ratios above the cap. This is not a phenomenon specific to Norway: Arena, Chen, Choi, Geng, Gueye, Lybek, Papageorgiou, and Zhang (2020) document substantial shares of mortgages with LTV ratios in excess of the cap across a broad set of advanced economies.

2.2 Economic downturn

During the fall of 2014, the international oil price collapsed from around $120 to less than $50. This triggered a local recession in southwestern Norway, where Norwegian oil production is concentrated. We define the recession area as Rogaland County, which at the time was one of 19 counties, containing about 10% of the total population.

Unemployment in Rogaland spiked in 2015, rose further in 2016, and began to decline in 2017 (Figure 1). Norwegian unemployment insurance is relatively generous, with a statutory replacement

ratio of 0.62, but the benefit is capped, so the effective replacement ratio for high-income workers is lower. Juelsrud and Wold (2025) estimate that workers laid off during the oil price recession faced an average replacement rate close to 0.50, comparable to the range of 0.30–0.50 across US states.[1]

Despite the sharp rise in unemployment, the aggregate income decline was moderate: total household income in Rogaland fell 2.3% from peak to trough, partly because the unemployment *rate* remained below 4% even at its peak.[2] However, Juelsrud and Wold (2025) show that even workers who remained employed but faced elevated job-loss *risk* increased their saving substantially, suggesting larger consumption effects than the moderate income decline would imply. House prices in Rogaland fell by eight percent from peak to trough. The downturn should therefore be viewed as a joint shock to expected income and asset prices.

While the oil price collapse had some negative spillovers outside Rogaland, the rest of the country did not experience a recession. Oslo, for instance, saw no increase in unemployment and no decline in house prices. In our empirical analysis, we document large consumption declines in Rogaland *relative* to the rest of the country. Although the recession was regional, Norges Bank responded by lowering the policy rate: the first cut came in December 2014, followed by additional cuts in June 2015, September 2015, and March 2016. These rate cuts provide the monetary policy variation we exploit in Section 7.

[1] See the Peter G. Peterson Foundation report at https://www.pgpf.org/article/budget-basics-unemployment-insurance-explained/.

[2] In 2016, there were 11,706 unemployed individuals in Rogaland, accounting for 3.9% of the population aged 18–66.

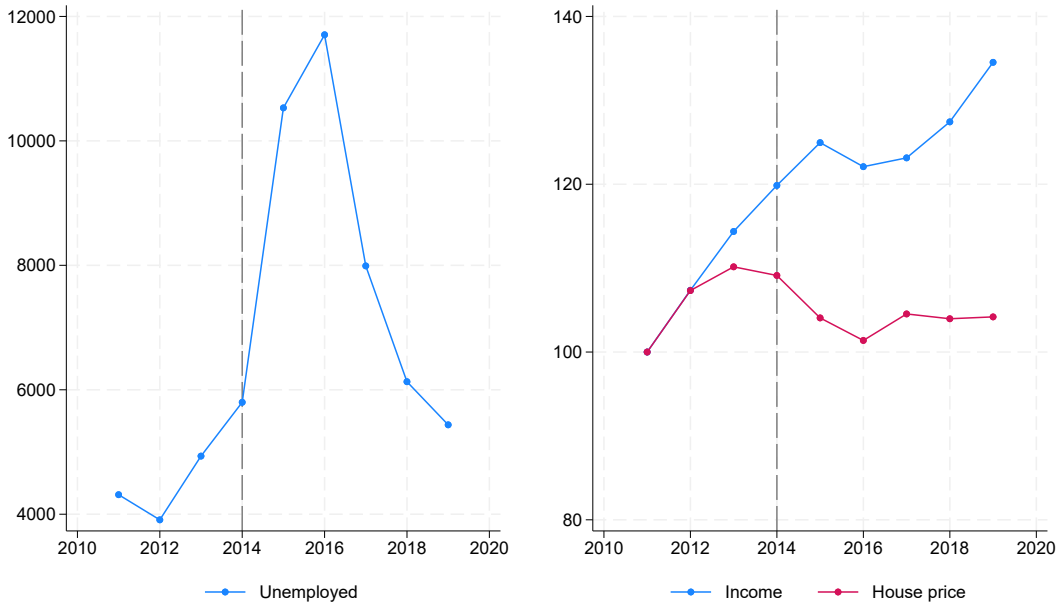


Figure 1: Macroeconomic developments in Rogaland.

Notes: The left panel shows total household income (2011=100) and the house price index for single-family homes (2011=100) in Rogaland county. The right panel shows the number of unemployed individuals in Rogaland county.

3 Data

We combine three Norwegian administrative data sets covering income and wealth, housing transactions, and consumption. Together, these data allow us to identify which households are affected by tighter loan-to-value caps, observe their balance sheet responses, and track their consumption through the subsequent economic downturn.

Income and wealth data Income and wealth data come from third-party reported tax filings covering the full population of Norwegian individuals. We observe liquid wealth (assets that can be quickly converted to cash), which is crucial for identifying potential home buyers, as well as income, total wealth, and parental wealth, which serve as key control variables. We identify mortgage debt from account-level data, restricting attention to accounts likely to represent mortgages based on debt amount and bank type. This approach allows us to construct more precise loan-to-value ratios than would be possible using total debt alone. All income and wealth variables are measured at year-end.

Housing transactions data We merge the tax data with housing transaction records from the Land Registry, which provide the exact purchase price for each transaction. Combined with our mortgage

debt measure, this allows us to compute loan-to-value ratios for all households that purchase a home in a given year. The transaction records also provide a purchase indicator that we use to identify home buyers and to classify households along the extensive margin.

Consumption data Our main outcome variable is household consumption, measured through debit card transactions. Norway is an almost cashless economy (Galaasen, Mæhlum, and Ahn, 2024), making card transactions a near-comprehensive measure of household spending. The data are provided by NETS, a major payment services provider in Scandinavia, and cover all debit card payments through the BankAxept system for the universe of Norwegian residents. During our sample period, BankAxept accounts for approximately 80% of all card transactions in Norway.[3] Transactions are classified into 26 consumption categories based on the COICOP system.[4] We aggregate the data to annual frequency for our main analysis. All data are collapsed to the household level, which is our unit of observation. Monetary values are expressed in USD using a constant exchange rate of 8 NOK per USD.

4 Identification strategy

Our identification strategy compares two cohorts of prospective first-time home buyers: those on the margin of purchasing in 2011, who faced a 90% LTV cap, and those on the margin of purchasing in 2012, who faced the tighter 85% cap. We focus on first-time buyers because they are the group most likely to face a binding LTV constraint.[5] By studying prospective buyers rather than actual buyers, our design captures both the intensive margin (balance sheet adjustments conditional on buying) and the extensive margin (changes in the probability of purchasing).

Treatment definition To identify prospective buyers, we construct a household-specific "wanted house price" for household i in year t, defined as the average transaction price in year t among buyers of similar age ($\pm$ 3 years), household size, and municipality. We then compute the ratio of each household's liquid wealth to its wanted house price. A household is classified as a prospective buyer in year t if it (a) does not own a home at the end of year $t-1$, (b) has liquid wealth at least equal to 10% of its wanted house price at the end of year $t-1$, and (c) did not satisfy condition (b) at the end of year $t-2$.

The 10% threshold corresponds to the minimum down payment under the pre-reform 90% LTV cap. Figure A.1 confirms that this threshold is empirically relevant: the probability of purchasing

[3] The aggregate debit card transactions data closely tracks national accounts data for household consumption and serves as a reliable real-time indicator. See Aastveit, Fastbø, Granziera, Paulsen, and Torstensen (2024).

[4] COICOP (Classification of Individual Consumption by Purpose) is developed by the United Nations Statistics Division and classifies expenditures by purpose, including categories such as food, clothing, housing, utilities, transport, and recreation (Eurostat 2023).

[5] Repeat buyers typically have equity from a previous property, making the LTV cap less likely to bind.

a home increases steeply in liquid wealth relative to the wanted house price up to 10%, and stabilizes thereafter. Condition (c) ensures that treatment and control households do not overlap, by restricting attention to households that newly crossed the wealth threshold in the relevant year.

Our treatment and control groups are thus defined as follows:

$$T_i = \begin{cases} 1, & \text{if household } i \text{ is a potential buyer in 2012, i.e. satisfies all of the following:} \\ & (a) \text{ was not a homeowner in 2011,} \\ & (b) \text{ had liquid wealth } \geq 0.10 \times \text{wanted house price at end of 2011,} \\ & (c) \text{ had liquid wealth } < 0.10 \times \text{wanted house price at end of 2010;} \\ 0, & \text{if household } i \text{ is a potential buyer in 2011, i.e. satisfies all of the following:} \\ & (a) \text{ was not a homeowner in 2010,} \\ & (b) \text{ had liquid wealth } \geq 0.10 \times \text{wanted house price at end of 2010,} \\ & (c) \text{ had liquid wealth } < 0.10 \times \text{wanted house price at end of 2009;} \end{cases}$$

Identifying assumptions Identification proceeds in two stages, each requiring its own assumptions. The first stage concerns the impact of the LTV tightening on household balance sheets. Here we must assume that households identified as prospective buyers in 2011 and 2012 would have had similar balance sheet outcomes absent the policy change.

We see two main threats to this assumption. First, because identification relies on time-series variation, macroeconomic shocks or trends between 2011 and 2012 could confound the comparison. However, the LTV distributions for buyers in 2011 and 2012 display clear bunching at exactly the prevailing cap in each year — 90% in 2011 and 85% in 2012 (Figure 2). It is highly unlikely that this pattern is driven by macroeconomic factors other than the LTV regulation itself.

Second, anticipation effects could affect household saving behavior before the new cap took effect. If households adjusted their savings in response to the announcement, those identified as prospective buyers in 2012 might differ systematically from those in 2011. However, the tightening was announced in December 2011, giving households very little time to adjust. Moreover, our classification relies on total liquid wealth — a stock variable — which is unlikely to change meaningfully in the few weeks between the announcement and the year-end measurement date.

The second stage concerns consumption during the downturn. For the recession results to be valid, treated and control households must have had similar consumption trajectories during the downturn absent the LTV tightening. Our dynamic triple-difference estimates confirm that consumption evolves on parallel trends prior to the recession (Figure 4), though we cannot rule out that the two groups would have *reacted* to the downturn differently for reasons unrelated to the

LTV policy.

The most natural concern is differential exposure to income shocks during the downturn. If treated households were less exposed to the oil sector or more exposed to the public sector, this could explain their smaller consumption declines. Table 1 shows that this is not the case: treated and control households have identical income shares from oil-related sectors (6.2%), and treated households are slightly *less* exposed to the public sector. Moreover, treated households have a *higher* probability of receiving unemployment insurance during the downturn (21% vs. 20%, $p < 0.01$). If anything, treated households face higher ex-ante and ex-post income risk, which would bias against finding a positive treatment effect on consumption. As a robustness exercise, we control explicitly for sectoral income shares in the analysis.

	T = 0	**T = 1**	**Difference (T=0 - T=1)**
Income share from oil (related) sectors	6.2	6.2	0.0 (0%)
Income share from public sector	33	32	1.0* (=3%)
Receives UI during post-period	20	21	1.4*** (=7%)

Table 1: Income risk by treatment status (%).

Notes: Income shares are calculated as share of total household income from given sectors using the SIC2007 industry classification. Oil-related sectors include Extraction of crude petroleum and natural gas (06) and Mining support service activities (09). Public sector includes Public administration and defence (84). Receives UI during post-period is the share of households receiving unemployment insurance during 2015–2017. Significance levels: $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5 Balance sheet effects of the LTV tightening

Figure 2 compares the distribution of estimated LTV ratios for prospective buyers who purchase in 2011 (control) versus 2012 (treatment). Two features stand out. First, LTV ratios are often high — frequently above the regulatory cap and sometimes exceeding 100%. As discussed in Section 2, this reflects banks' room to deviate from the cap, the use of parental collateral, and measurement error in mortgage debt.

Second, while there is a clear shift in the peak of the distribution from 0.90 to 0.85, the aggregate impact on LTV ratios is small. The effect is concentrated among mortgages with LTV ratios very close to the cap; the rest of the distribution is largely unaffected. Because these marginal mortgages account for a small share of the total, the overall effect on average LTV ratios among buyers is limited.

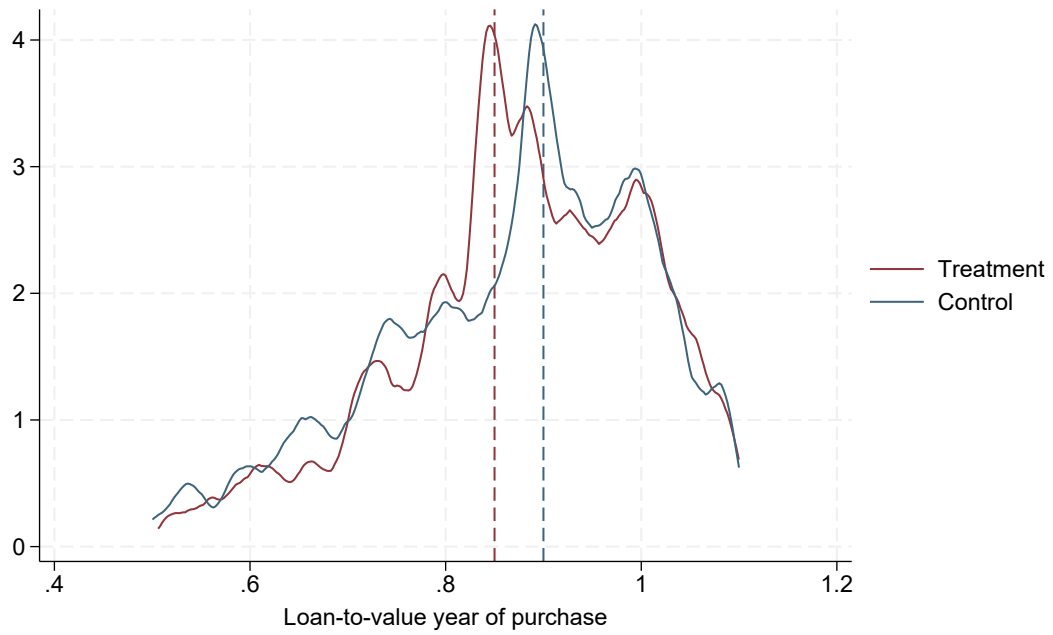


Figure 2: Loan-to-value for treatment and control households.

Notes: Estimated mortgage debt relative to house purchase price (LTV). House buyers in year of potential buyer. Treatment households buy in year 2012. Control households buy in year 2011. LTV∈[0.5,1.1].

Table 2 quantifies these effects. In columns 1 and 2, we regress LTV ratios on treatment status among households that purchase a home. Our LTV measure uses purchase price in the denominator and is therefore defined only for buyers. Considering all buyers with LTV ratios in [0.50, 1.10] — the sample shown in Figure 2 — we find no significant effect of treatment on LTV ratios (column 1). Restricting the sample to LTV ratios in [0.75, 1.00], we find a small but significant reduction of 0.6 percentage points (column 2). The intensive margin effect on leverage conditional on purchase is thus visible but modest.

Column 3 reports the effect on the probability of purchasing a home. Treated households have a 4.0 percentage point lower purchase probability, a reduction of nearly 28% relative to the control group mean of 14.5%. Because first-time home purchases are associated with large debt increases, this has substantial implications for total debt. Column 4 confirms this: treated households accumulate $73,560 less debt. When we add a control for whether the household actually purchases (column 5), the coefficient on treatment becomes small and insignificant, suggesting that the debt reduction operates primarily through the extensive margin — through canceled or postponed purchases — rather than through lower borrowing among those who still buy.

	(1) LTV∈[0.5,1.1]	(2) LTV∈[0.75,1.0]	(3) Buy	(4) ΔDebt	(5) ΔDebt
T	0.007	-0.006*	-0.040***	-73,560***	-15,920
	(0.005)	(0.003)	(0.004)	(11,482)	(10,244)
Buy					1,434,782***
					(15,477)
Constant	0.868***	0.889***	0.145***	343,788***	135,570***
	(0.003)	(0.002)	(0.003)	(8087)	(7544)
Observations	2,713	1,728	32,925	32,925	32,925

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2: Balance sheet effects of LTV-tightening

Notes: Estimated mortgage debt relative to purchase price (LTV), indicator variable for house purchase (Buy) and change in estimated mortgage debt (ΔDebt) – in year of potential buyer. $T_i = 1$ for treated households, i.e. potential house buyers in 2012. $T_i = 0$ for control households, i.e. potential house buyers in 2011.

Appendix Table B.1 provides complementary evidence on purchase prices and financial wealth. Conditional on buying, households subject to the tighter LTV cap purchase less expensive homes. Financial wealth increases for the treatment group overall, but not when conditioning on purchase. This pattern is consistent with two offsetting effects: buyers face a negative liquidity effect (higher down payment requirements deplete liquid savings), while non-buyers retain their savings, generating a positive liquidity effect through the extensive margin.

5.1 Balance sheet differences at the time of the downturn

The LTV tightening occurred in 2012, but the recession we study began in late 2014. For our consumption analysis to be informative, the policy-induced balance sheet differences must persist over this two-year interval. We verify that this is the case.

Table 3 reports balance sheet characteristics by treatment status in 2014. Treated households have 17% lower housing wealth and are 4 percentage points (11%) less likely to be homeowners. They carry 8% less debt, and their debt-to-income and debt-to-wealth ratios are 5.3% and 3.3% lower, respectively. Financial wealth is slightly higher for the treatment group, but the difference is small in both statistical and economic terms. Overall, the policy-induced differences in leverage remain large and economically meaningful at the onset of the recession.

	T = 1	T = 0	Difference (T=1 − T=0)
Primary housing wealth	992,251	1,191,796	-199,545*** (-17.0%)
Home owner	0.37	0.41	-0.04*** (-11.0%)
Debt	1,510,226	1,641,291	-131,065*** (-8.0%)
Debt / Income	2.51	2.64	-0.13*** (-5.3%)
Debt / Gross wealth	2.10	2.20	-0.07* (-3.3%)
Financial wealth	304,598	297,492	7,106* (2.4%)

Table 3: Summary statistics by treatment status in 2014

Notes: The table reports means by treatment status. The difference column reports $T = 1 - T = 0$, with percentage differences computed relative to the $T = 0$ mean. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

6 The impact of LTV caps on household consumption during a downturn

6.1 LTV caps and average consumption

We now study how the LTV tightening affects household consumption during the downturn. We normalize consumption by each household's pre-recession level, defining $\overline{C}_{i,t} = C_{i,t}/C_{i,2014}$, and estimate the following triple difference-in-differences specification:

$$\overline{C}_{it} = \alpha_i + \delta_t + \Gamma_1 \left(Rogaland_i \times \delta_t \right) + \Gamma_2 \left(T_i \times \delta_t \right) + \Gamma_3 \left(Rogaland_i \times T_i \times \delta_t \right) + \varepsilon_{it} \tag{1}$$

where α_i and δ_t are household and year fixed effects, $Rogaland_i$ indicates residence in the recession region, and T_i indicates treatment status. The coefficients of interest are the Γ_3 vector, which captures how consumption evolves for treated households relative to control households within Rogaland, year by year. In some specifications we add controls for age, income, wealth, and parental wealth (measured in 2014), each interacted with year indicators and the Rogaland dummy. We also verify robustness to controlling for sectoral income shares. Standard errors are clustered at the household and year-municipality level.

We begin by documenting the recession itself. Figure 3 plots the estimated Γ_1 coefficients, which capture the evolution of consumption in Rogaland relative to the rest of Norway. Prior to the downturn, consumption in Rogaland tracks the rest of the country closely. Starting in 2015, consumption in Rogaland declines sharply. By 2017, relative consumption is nearly 11% lower.

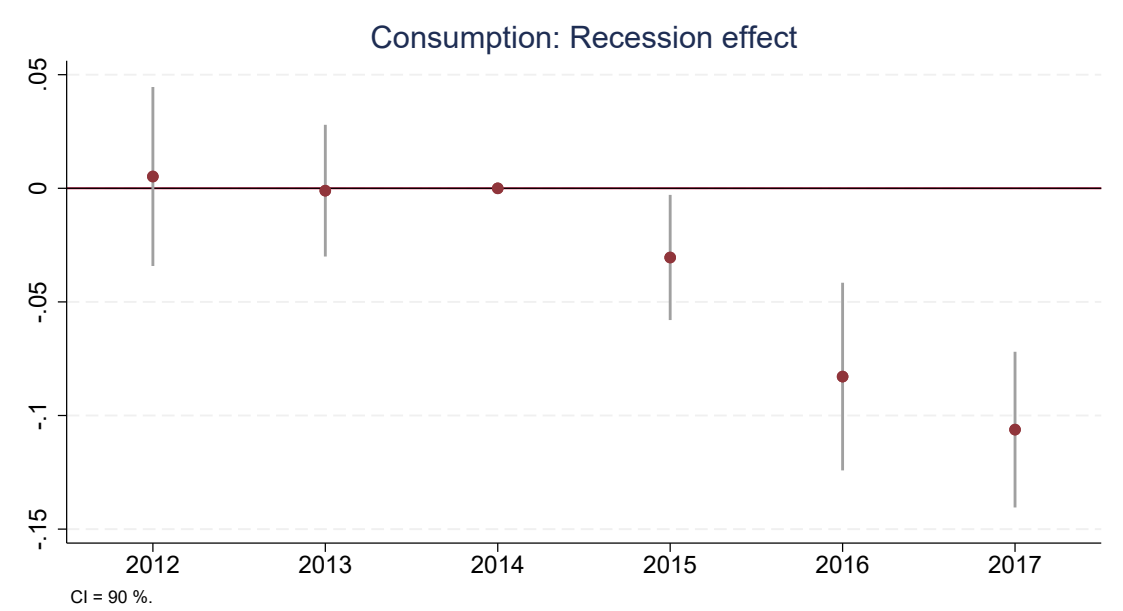


Figure 3: Recession effect: Estimated Γ_1-coefficients from equation (1).

Notes: Consumption dynamics in Rogaland. Consumption is defined as $\overline{C}_{it} = C_{i,t}/C_{i,2014}$. $T_i = 1$ for treated households, i.e. potential house buyers in 2012. $T_i = 0$ for control households, i.e. potential house buyers in 2011. No control variables.

Figure 4 plots the Γ_3 coefficients — the differential consumption effect for treated households in Rogaland. Prior to the recession, treated and control households display parallel consumption trends, supporting our identification strategy. After 2014, treated households experience substantially smaller consumption declines. By 2017, the treatment effect reaches approximately 9% in the specification with balance sheet controls, implying that the recession-induced consumption decline is largely offset for this group.

The results are qualitatively similar without control variables (Figure A.2), though the magnitudes are somewhat smaller: the treatment effect peaks at 7.5% in 2016 and is approximately 5% in 2017.

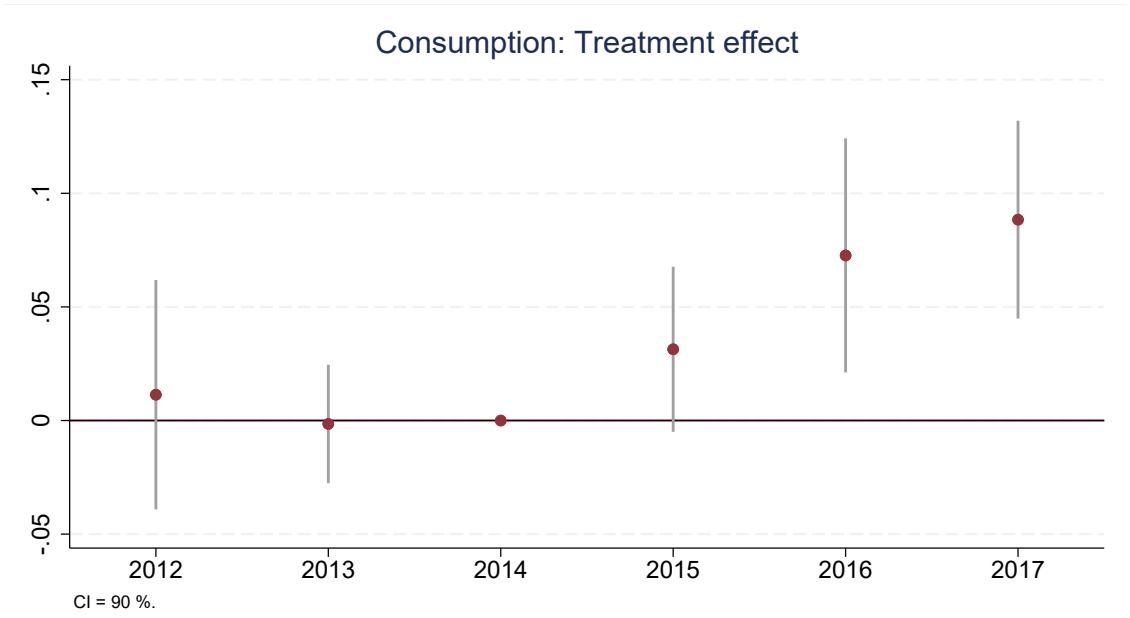


Figure 4: Treatment effect: Estimated Γ_3-coefficients from equation (1)

Notes: Consumption dynamics in Rogaland for treated households. Consumption is defined as $\overline{C}_{it} = C_{i,t}/C_{i,2014}$. $T_i = 1$ for treated households, i.e. potential house buyers in 2012. $T_i = 0$ for control households, i.e. potential house buyers in 2011. Control variables include 2014-values of age, income, wealth and parental wealth, interacted with year and Rogaland-dummy.

Having confirmed parallel pre-trends, we collapse the estimation to a pre-versus-post comparison. We define $Post_t = 1$ if year ≥ 2015 and estimate:

$$\overline{C}_{it} = \alpha_i + \delta_t + \beta_1 \left(Rogaland_i \times Post_t\right) + \beta_2 \left(T_i \times Post_t\right) + \beta_3 \left(Rogaland_i \times T_i \times Post_t\right) + \varepsilon_{it} \quad (2)$$

The recession effect is captured by β_1 and the treatment effect by β_3.

Table 4 reports the results. The recession coefficient (β_1) indicates that consumption in Rogaland falls 7.5% relative to the rest of Norway during 2015–2017 (columns 1–2). Columns 3–5 include municipality-by-year fixed effects, which absorb this coefficient.

The treatment effect (β_3) is positive and significant across all specifications. In column 1, treated households have 5.0% higher relative consumption during the recession. Adding household fixed effects (column 2) increases the estimate to 5.5%. With municipality-by-year fixed effects (column 3), the treatment effect is 6.2%, and with the full set of balance sheet and sectoral income controls (column 5), it reaches 7.1%. The consumption stabilization effect of tighter LTV regulation is thus robust to a demanding set of fixed effects and controls.

	(1) C	(2) C	(3) C	(4) C	(5) C
Post×Rogaland	-0.075***	-0.075***			
	(0.017)	(0.017)			
Post×T	0.005	0.005	0.004	-0.009	-0.009
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Rogaland×T	-0.008	-0.120*	-0.151**	-0.132*	-0.153*
	(0.014)	(0.069)	(0.069)	(0.079)	(0.079)
Post×Rogaland×T	0.050**	0.055**	0.062***	0.071***	0.071***
	(0.023)	(0.023)	(0.023)	(0.025)	(0.026)
Year FE	Yes	Yes	-	-	-
HH FE	No	Yes	Yes	Yes	Yes
Muni*Year FE	No	No	Yes	Yes	Yes
Balance sheet controls	No	No	No	Yes	Yes
Sectoral income controls	No	No	No	No	Yes
R2	0.032	0.415	0.428	0.436	0.433
N	177,247	177,247	177,088	137,392	122,062

Standard errors in parentheses, clustered at household and year-municipality level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Regression results from estimating equation (2).

Notes: Consumption is defined as $\overline{C}_{it} = C_{i,t}/C_{i,2014}$. Post=1 if year≥2015, and zero otherwise. $T_i = 1$ for treated households, i.e. potential house buyers in 2012. $T_i = 0$ for control households, i.e. potential house buyers in 2011. Balance sheet control variables include 2014-values of age, income, wealth and parental wealth, interacted with year and Rogaland-dummy. Sectoral income control variables include 2014-values of income share from oil related sectors and public sector, interacted with year and Rogaland-dummy.

6.2 Which consumption groups are affected?

We decompose total consumption into six groups based on durability, housing linkage, and cyclicality, as summarized in Table 5. The key distinctions are between non-durable necessities (food and services, which together account for over half of total spending and display limited cyclicality), housing-related durables (utilities and housing, which include furniture, equipment, and maintenance), and transport (dominated by vehicle purchases). The remaining categories — personal recreation and personal goods — capture intermediate forms of discretionary spending.

Consumption group	**COICOP2018 group**	**Content**
Services (33%)	06, 07.3–4, 08, 09.4, 10, 11, 12, 12.6, 12.7, 13.3, 14	Health; transportation services; information and communication; recreational services; education services; restaurants and accommodation services; insurance and financial services; social protection
Food (21%)	01, 02	Food and non-alcoholic beverages; alcoholic beverages; tobacco and narcotics
Utilities and housing (19%)	04, 05	Housing; water; electricity; gas and other fuels; furnishing; household equipment; routine household maintenance
Personal recreation (11%)	09.1–3, 09.5	Recreational durables; other recreational goods; garden products and pets; cultural goods
Transport (10%)	07.1–2	Purchase of vehicles; operation of personal transport equipment
Personal goods (5.3%)	03, 13.1, 13.2	Clothing and footwear; personal care; other personal effects

Table 5: Consumption group classification

Notes: Consumption is divided into six distinct groups. Their share of total consumption in 2014 is reported in parenthesis in the first column.

We re-estimate equation (2) separately for each consumption group. Figure 5 plots the recession effect ($\hat{\beta}_1$) against the treatment effect ($\hat{\beta}_3$) for each group. The figure reveals a striking pattern: consumption groups with the largest recession-induced declines also exhibit the largest treatment effects. Transport and utilities — which fall most sharply during the recession — show the largest attenuation for treated households. Food and services show minimal recession sensitivity and correspondingly small, insignificant treatment effects. The correlation between recession effects and treatment effects across groups is 0.85, providing within-household validation that the consumption stabilization occurs precisely where the downturn hits hardest.

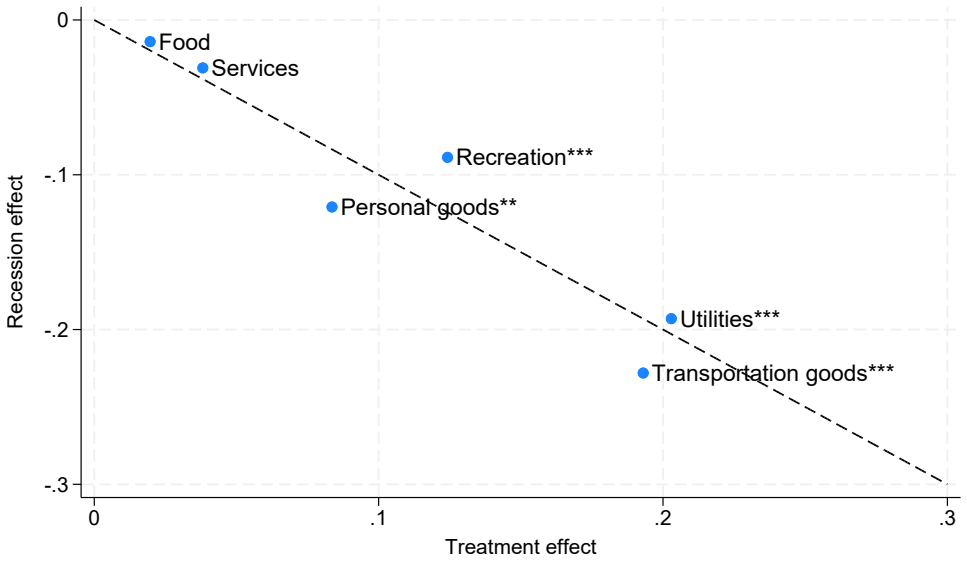


Figure 5: Recession effect vs. treatment effect by consumption group

Notes: Recession effect is $\hat{\beta}_1$ from equation (2). Treatment effect is $\hat{\beta}_3$ from equation (2). Equation (2) is estimated separately for each of the six consumption groups. Stars next to consumption group name signify significance.

6.3 Decomposing the consumption effect

Our main finding is that the consumption fall during the recession is less severe for households exposed to the tighter LTV cap. Intuitively, this effect can be driven by households who are deterred from entering the housing market, or by households who still buy a house but adjust their balance sheets in response to the reform. One highly relevant adjustment margin is lower leverage, but it could also include changes in liquidity, interest rate payments, and similar channels.

Separating the extensive margin from the intensive margin is not straightforward, as these channels are not separately identified without further assumptions. In this section we provide more structure, which allows us to point-identify some components of the total effect and put bounds on the remaining components. Under the maintained assumptions, the extensive margin can account for essentially the entire treatment effect, while the intensive margin is quantitatively small. This means that LTV caps are effective in stabilizing consumption during a downturn, but only insofar as they keep households from becoming homeowners.

6.3.1 Structural framework

Consider two policy regimes, indexed by $d \in \{0, 1\}$. Regime $d = 0$ corresponds to a policy regime with no LTV-caps, while regime $d = 1$ corresponds to a tighter borrowing environment resulting with LTV-caps in place. Let $Y(d)$ denote a household's consumption change during the recession under policy regime d. The triple difference-in-differences estimator identifies the average total

effect of moving from the loose to the tight regime,

$$TE \equiv E[Y(1) - Y(0)]\,.$$

Households may respond to the policy change by adjusting their home-purchase decisions. Let $M(d) \in \{0, 1\}$ denote the household's decision to purchase a home under policy regime d, where $M(d) = 1$ indicates purchasing and $M(d) = 0$ indicates renting. Define household type by the pair of potential purchase decisions,

$$S \equiv (M(0), M(1)).$$

This classification partitions households into four principal types:

- Always buyers: $S = (1, 1)$,
- Always renters: $S = (0, 0)$,
- Switchers: $S = (1, 0)$,
- Defiers: $S = (0, 1)$.

Using this partition, the total treatment effect can be written as a weighted average of type-specific treatment effects

$$TE = \sum_{s \in \{(1,1),(0,0),(1,0),(0,1)\}} \Pr(S = s)\, E[Y(1) - Y(0) \mid S = s]\,. \tag{3}$$

We are mainly interested in identifying the treatment effects for the always-buyers, which we refer to as the intensive margin, and the switchers, which we refer to as the extensive margin.

6.3.2 Assumptions

Decomposing the treatment effect into the respective contributions from the different household types requires additional assumptions. Here we make three assumptions in order to bound the intensive and extensive margin effects.

Monotonicity We impose a monotonicity restriction on the home-purchase decision,

$$M(1) \leq M(0),$$

which rules out the possibility that introducing an LTV cap induces households to purchase a home. This means that no households are "defier" households, so that $Pr\big(S = (0, 1)\big) = 0$.

Further, monotonicity means that all households that purchase in the post-tightening period must be always buyers, implying

$$\Pr\big(S = (1, 1)\big) = \pi_1,$$

where $\pi_1 \equiv \Pr(M = 1 \mid d = 1)$ denotes the observed post-tightening purchase probability. Similarly, all households that rent in the pre-tightening period must be always renters, implying

$$\Pr\big(S = (0,0)\big) = 1 - \pi_0,$$

where $\pi_0 \equiv \Pr(M = 1 \mid d = 0)$ denotes the observed pre-tightening purchase probability. The remaining households – those who purchase under the loose regime but not under the tight regime – are switchers, with population share

$$\Pr\big(S = (1,0)\big) = \pi_0 - \pi_1.$$

Let

$$\tau_{ij} \equiv E[Y(1) - Y(0) \mid S = (i,j)]$$

denote the average treatment effect within stratum $S = (i,j)$. Under monotonicity, the total treatment effect from equation (3) can be expressed as

$$TE = \pi_1\, \tau_{1,1} + (1 - \pi_0)\, \tau_{0,0} + (\pi_0 - \pi_1)\, \tau_{1,0}. \tag{4}$$

No effect on always renters We further impose the restriction that the policy tightening has no effect on households that rent under both regimes (always renters). Given that our empirical estimates should be viewed as partial equilibrium effects, we view this assumption as reasonable. The implication is, of course, that

$$\tau_{0,0} \equiv E[Y(1) - Y(0) \mid S = (0,0)] = 0.$$

Under this assumption, the total treatment effect from equation (4) simplifies further to

$$TE = \pi_1\, \tau_{1,1} + (\pi_0 - \pi_1)\, \tau_{1,0}. \tag{5}$$

Resilience ordering. We view the two assumptions made so far, monotonicity and no effect on always-renters, as relatively innocuous. However, as seen from equation (5), these assumptions are not sufficient to separately identify the intensive and extensive margins. Conceptually, the treatment effect on always-buyers and switchers are not identified because these two groups cannot be separated under the loose policy regime.

To bound the intensive and extensive margin effects we make one additional assumption. Aastveit et al. (2024) study the same LTV caps as in this paper and show that the households which are deterred from buying have weaker balance sheets than non-deterred households. We therefore assume that, under the loose regime – when *both groups are buyers* – always-buyers have

consumption changes that are at least as large as switchers. Or in other words, if both groups experience consumption declines during the recession, the decline is at least as big for switchers as for always-buyers. Formally,

$$E[Y(0) \mid S = (1,1)] \;\geq\; E[Y(0) \mid S = (1,0)].$$

Let

$$a_0 \equiv E[Y(0) \mid S = (1,1)], \qquad s_0 \equiv E[Y(0) \mid S = (1,0)].$$

The observed mean outcome among pre-reform buyers satisfies

$$\mu_{0,1} = \frac{\pi_1}{\pi_0} a_0 + \frac{\pi_0 - \pi_1}{\pi_0} s_0.$$

Because $\mu_{0,1}$ is a convex combination of a_0 and s_0, the resilience ordering implies

$$s_0 \;\leq\; \mu_{0,1} \;\leq\; a_0.$$

6.3.3 Bounding the extensive and intensive margins

We now have the assumptions we need to say something quantitatively meaningful about the intensive vs extensive margin effects.

The intensive margin Recall that

$$\tau_{1,1} = E[Y(1) \mid S = (1,1)] - E[Y(0) \mid S = (1,1)] = \mu_{1,1} - a_0,$$

Using the resilience ordering, this implies

$$\tau_{1,1} \;\leq\; \mu_{1,1} - \mu_{0,1}. \tag{6}$$

Equation (6) states that the intensive-margin effect is bounded above by the difference in observed outcomes for buyers across the two policy regimes. The intuition is that, under the loose regime, the group of buyers consists of both always-buyers and switchers. By the resilience ordering, always-buyers have (weakly) higher baseline consumption changes than switchers, implying that the average outcome among pre-reform buyers, $\mu_{0,1}$, is weakly lower than the counterfactual outcome for always-buyers alone. As a result, $\mu_{1,1} - \mu_{0,1}$ overstates the true treatment effect for always-buyers, yielding an upper bound on $\tau_{1,1}$.

The extensive margin Recall that

$$\tau_{1,0} = E[Y(1) \mid S = (1,0)] - E[Y(0) \mid S = (1,0)].$$

Using the resilience ordering, we obtain

$$\tau_{1,0} \;\geq\; E[Y(1) \mid S = (1,0)] - \mu_{0,1}. \tag{7}$$

Equation (7) states that the treatment effect for switchers is bounded below by the difference between their post-reform outcome and the observed pre-reform mean among buyers. The intuition is that, under the loose regime, buyers consist of both always-buyers and switchers. By the resilience ordering, always-buyers have (weakly) higher baseline outcomes than switchers, implying that the average outcome among pre-reform buyers, $\mu_{0,1}$, is weakly higher than the counterfactual outcome for switchers alone. As a result, subtracting $\mu_{0,1}$ from $E[Y(1) \mid S = (1,0)]$ understates the true treatment effect for switchers, yielding a lower bound on $\tau_{1,0}$.

To obtain a measure for $E[Y(1) \mid S = (1,0)]$, note that those who do not purchase in the post-period are a mix of always-renters and switchers. Hence,

$$\mu_{1,0} \equiv E[Y \mid d = 1, M = 0] = \frac{1-\pi_0}{1-\pi_1} E[Y(1) \mid S = (0,0)] + \frac{\pi_0 - \pi_1}{1-\pi_1} E[Y(1) \mid S = (1,0)].$$

Solving for $E[Y(1) \mid S = (1,0)]$ yields

$$E[Y(1) \mid S = (1,0)] = \frac{(1-\pi_1)\mu_{1,0} - (1-\pi_0)E[Y(1) \mid S = (0,0)]}{\pi_0 - \pi_1}.$$

Using the assumption of no effect on always-renters, $E[Y(1) \mid S = (0,0)] = E[Y(0) \mid S = (0,0)] = \mu_{0,0}$, we obtain

$$E[Y(1) \mid S = (1,0)] = \frac{(1-\pi_1)\mu_{1,0} - (1-\pi_0)\mu_{0,0}}{\pi_0 - \pi_1}.$$

6.3.4 Quantitative take-away

Table 6 lists the relevant data moments. The total treatment effect TE is the coefficient on the triple interaction term, equal to 0.05. This is the mitigating consumption effect from Table 4. π_0 and π_1 capture the purchase probabilities under the different regimes, and are measured as the average purchase probability for the treatment and control group. $\pi_0 = 0.145$ means that 14.5% of control households purchase a house, while $\pi_1 = 0.105$ means that 10.5% of treatment households purchase a house. Together these moments imply that the population share for switchers is $\pi_0 - \pi_1 = 0.04$.

We also need to estimate $\mu_{0,0}$, $\mu_{0,1}$, $\mu_{1,0}$ and $\mu_{1,1}$. These are the average consumption outcomes for buyers and non-buyers under each regime. To do this we take our triple interaction specification and add an extra variable $Buyer_i$:

$$\overline{C}_{it} = \alpha_i + \delta_t + \beta_1(Rogaland_i \times Post_t) + \beta_2(T_i \times Post_t) + \beta_3(Buyer_i \times Post_t) + \beta_4(Rogaland_i \times T_i \times Post_t) + \beta_5(Rogaland_i \times Buyer_i \times Post_t) + \beta_6(T_i \times Buyer_i \times Post_t) + \beta_7(Rogaland_i \times T_i \times Post_t \times Buyer_i) + \varepsilon_{it}. \tag{8}$$

$Buyer_i = 1$ if the household purchases a house in the year of being a predicted buyer, and zero otherwise. How the estimated coefficients from equation (8) map into the μ's is specified in Table 6.

Model moment	Empirical counterpart (data moment)	Value
$TE \equiv E[Y(1) - Y(0)]$	Coefficient on triple diff from baseline specification	0.050
$\pi_0 \equiv \Pr(M(0) = 1)$	$\Pr(M = 1 \mid \texttt{TREAT} = 0)$	0.145
$\pi_1 \equiv \Pr(M(1) = 1)$	$\Pr(M = 1 \mid \texttt{TREAT} = 1)$	0.105
$\mu^{obs}_{0,0} \equiv E[Y \mid D = 0, M = 0]$	$\hat{\beta}_1$	-0.082
$\mu^{obs}_{0,1} \equiv E[Y \mid D = 0, M = 1]$	$\hat{\beta}_1 + \hat{\beta}_5$	-0.045
$\mu^{obs}_{1,0} \equiv E[Y \mid D = 1, M = 0]$	$\hat{\beta}_1 + \hat{\beta}_4$	-0.025
$\mu^{obs}_{1,1} \equiv E[Y \mid D = 1, M = 1]$	$\hat{\beta}_1 + \hat{\beta}_4 + \hat{\beta}_5 + \hat{\beta}_7$	-0.018
$E[Y(1) \mid S = (1,0)]$	$\dfrac{(1-\pi_1)\mu^{obs}_{1,0} - (1-\pi_0)\mu^{obs}_{0,0}}{\pi_0 - \pi_1}$	1.193

Table 6: Model moments and empirical counterparts used in the decomposition and bounds analysis

We now have what we need to quantify the bounds. From equation (6) we have that the treatment effect on always-buyers is bounded above, such that

$$\tau_{1,1} \;\leq\; \mu_{1,1} - \mu_{0,1} = -0.018 - (-0.045) = 0.027. \tag{9}$$

From equation (7) we have that the treatment effect on switchers is bounded below, such that

$$\tau_{1,0} \;\geq\; E[Y(1) \mid S = (1,0)] - \mu_{0,1} = 1.193 - (-0.045) = 1.238.$$

Recall that the total treatment effect is given by

$$TE = \pi_1 \, \tau_{1,1} + (\pi_0 - \pi_1) \, \tau_{1,0}.$$

The intensive-margin contribution, defined as $\pi_1\tau_{1,1}$, is *at most* equal to $0.105 \times 0.027 = 0.0028$, corresponding to approximately six percent of the total treatment effect. The extensive-margin contribution, defined as $(\pi_0 - \pi_1)\tau_{1,0}$, is *at least* equal to $0.04 \times 1.238 = 0.0495$, corresponding to

about 99 percent of the total treatment effect.[6]

The large magnitude of the extensive-margin component reflects that the share of switchers, $\pi_0 - \pi_1 = 0.04$, is small. Consequently, the expression for $E[Y(1) \mid S = (1,0)]$ involves dividing by a small quantity and is highly sensitive to small changes in the underlying moments. The quantitative magnitudes of the bounds should therefore be interpreted with caution. The robust conclusion from this exercise is qualitative: under the maintained assumptions, the extensive margin can account for essentially the entire treatment effect, while the intensive margin is bounded to play a quantitatively limited role.

7 The interaction of LTV caps and monetary policy

A key question for policymakers with both macroprudential and monetary policy tools is how they interact. The LTV tightening in 2012 left treated households with lower leverage when the 2014 recession hit. These balance sheet differences may affect the transmission of monetary policy during the downturn. Lower leverage could weaken the consumption response to monetary easing, as households have less debt whose servicing costs decline when rates fall. At the same time, lower leverage makes households more resilient, reducing the need for monetary stimulus in the first place. We study this interaction empirically.

7.1 Empirical approach

To analyze the effects of monetary policy on household consumption, we follow the local projections approach implemented in Norway by Holm et al. (2021). We also use the same narrative monetary policy shocks as in their paper and a closely aligned dependent variable – consumption growth. The shocks are derived from policy announcements and central bank communications, identifying the surprise component of policy decisions. A negative shock represents a surprise easing (lower rates than expected). As a first step, we verify in our data that these monetary policy shocks have the expected average effect on consumption.

Let q index quarters and let C_{iq} denote quarterly household consumption. For each horizon $h = 0, 1, \ldots, 8$, we estimate local projections of the form

$$\frac{C_{i,q+h} - C_{i,q-1}}{\overline{C}_{q-1}} = \alpha_i^h + \beta^h MP_q + \sum_{k=1}^{3} \gamma_k^h MP_{q-k} + \sum_{m=1}^{2} \rho_m^h \frac{C_{i,q-m} - C_{i,q-m-1}}{\overline{C}_{q-m-1}} + \eta_{\tau(q)} + \varepsilon_{iq}^h, \quad (10)$$

where MP_q is the narrative monetary policy shock from Holm et al. (2021), α_i^h is a household fixed effect, and $\eta_{\tau(q)}$ are quarter-of-year fixed effects. The variable $\overline{C}_{q-1}$ denotes average household consumption in quarter $q-1$, computed across all households in the sample. Hence, the dependent

[6] In principle, the intensive margin effect could be negative and the extensive margin effect could exceed 100%.

variable measures the cumulative change in household consumption from quarter $q-1$ to quarter $q+h$, scaled by lagged average consumption. As in Holm et al. (2021), we include three lags of the monetary policy shock and two lags of consumption growth as controls.

This specification is very close to the one in Holm et al. (2021). The main difference is mechanical and comes from the fact that our consumption data are observed directly from debit card transactions at quarterly frequency, while their consumption measure is constructed from administrative data at annual frequency. Conceptually, however, the exercise is the same: we estimate dynamic responses of household consumption to identified monetary policy shocks using local projections. In particular, our normalization by lagged average consumption follows their approach and makes the coefficients directly comparable to aggregate responses. By contrast, normalizing by each household's own lagged consumption would produce the average of household-level percentage changes, which is more sensitive to extreme observations and is not directly comparable to aggregate impulse responses.

Figure 6 reports the estimated coefficients β^h from equation (10). The figure shows that expansionary monetary policy shocks increase consumption on average, as expected. The response is negative already on impact, becomes larger over the subsequent quarters, and reaches its maximum within the first year, peaking around the third quarter. Thereafter the response remains negative but becomes less precisely estimated. This pattern suggests that the main effect of monetary policy on consumption occurs within roughly one year.

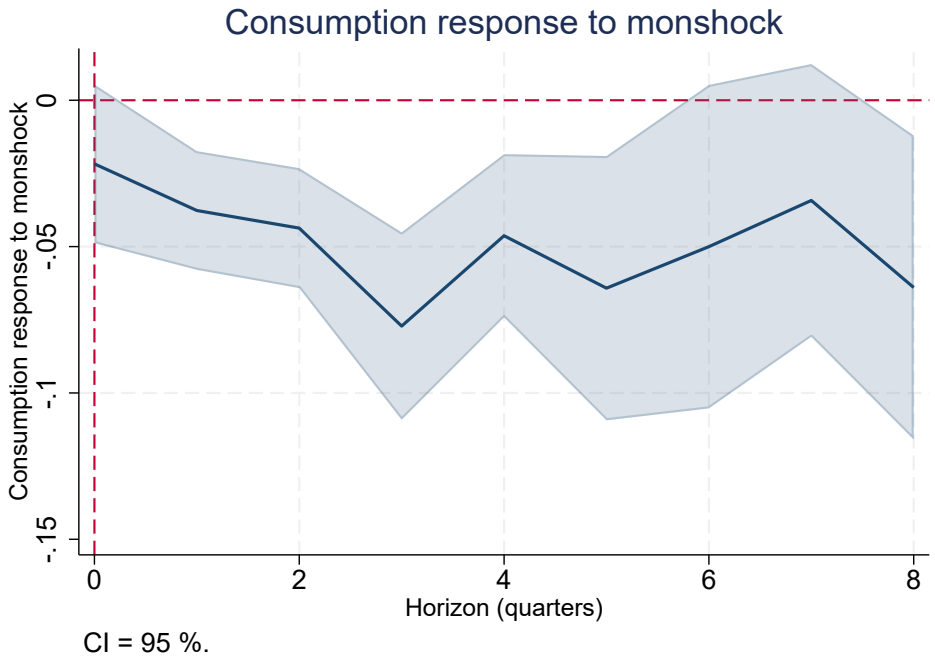


Figure 6: Consumption response to narrative monetary policy shocks: local projections

Notes: The figure plots the estimated coefficients β^h from equation (10) for horizons $h = 0, \ldots, 8$, together with 95% confidence intervals. The dependent variable is the cumulative change in quarterly household consumption from quarter $q-1$ to quarter $q+h$, scaled by average household consumption in quarter $q-1$. Monetary policy shocks are the narrative shocks from Holm et al. (2021).

7.2 Empirical specification

Motivated by the local projection results above, we construct a forward-looking consumption growth measure to study the interaction between macroprudential policy and monetary policy. Specifically, for household i in quarter q we define

$$\widetilde{C}_{iq} = \frac{C_{iq}^{\text{next year}} - C_{iq}^{\text{past year}}}{\overline{C}_q^{\text{past year}}}, \tag{11}$$

where

$$C_{iq}^{\text{past year}} = C_{i,q-1} + C_{i,q-2} + C_{i,q-3} + C_{i,q-4},$$

$$C_{iq}^{\text{next year}} = C_{i,q+1} + C_{i,q+2} + C_{i,q+3} + C_{i,q+4},$$

and $\overline{C}_q^{\text{past year}}$ is the cross-sectional average of $C_{iq}^{\text{past year}}$ in quarter q. Thus, this variable measures consumption over the following year relative to consumption over the previous year, normalized by lagged average consumption across households.

We use this measure as a dependent variable, which has three advantages in our setting. First, it

captures the lagged transmission of monetary policy over the horizon where Figure 6 indicates that most of the action occurs, namely within the first year. Second, by summing over four quarters on both sides, it addresses seasonal patterns in household consumption. Third, it is closely aligned with the empirical strategy in Holm et al. (2021), in the sense that we study cumulative consumption responses scaled by average lagged consumption.

We then estimate the following specification:

$$\widetilde{C}_{iq} = \alpha_i + \delta_q + \beta_1 \left(Post_q \times Rogaland_i \times T_i \times MP_q \right) + \Gamma X_{iq} + u_{iq}, \tag{12}$$

where α_i denotes household fixed effects and δ_q denotes year-quarter fixed effects in the specifications that include them. $Post_q$ equals one in the recession period and zero otherwise. With quarterly data, we define the post period to start in 2014Q4, when Norges Bank started cutting rates in response to the oil-price shock. Because the dependent variable is forward-looking, quarter 2014Q3 is also assigned to the post period, since its forward window already includes 2014Q4. $Rogaland_i$ is an indicator for residing in the recession region – Rogaland, T_i is the treatment indicator for exposure to the tighter 2012 LTV cap, and MP_q is the monetary policy shock in quarter q.

Besides the narrative monetary policy shocks from Holm et al. (2021), we also use high-frequency shocks, which identify Norwegian monetary policy surprises using changes in market interest rates in narrow windows around policy announcements, as estimated by Brubakk et al. (2022). Both shock series are available at quarterly frequency and capture unexpected changes in monetary policy. A positive shock represents a surprise tightening (higher rates than expected), while a negative shock represents a surprise easing (lower rates than expected). For our purposes, we examine how household consumption responds to these surprises, and whether the response differs between treated and control households during the recession.

The vector X_{iq} contains all the four variables included in the quadruple interaction, all possible interactions among them, and the same household-level controls used elsewhere in the paper, now interacted flexibly with the post and regional dimensions and with the monetary policy shock. Specifically, we include age, log post-tax income, log gross wealth, and log parental wealth, as well as income shares from oil-related sectors and the public sector. Standard errors are two-way clustered at the household and year-quarter levels. In some specifications year-quarter fixed effects absorb the main effect of MP_q, so the coefficient on MP_q itself is not separately identified there.

The coefficient of primary interest is β_1, on the quadruple interaction $Post_q \times Rogaland_i \times T_i \times MP_q$. This coefficient measures whether the consumption response to monetary policy differs for households that were exposed to tighter macroprudential regulation and are located in the recession region during the recession period. A positive coefficient implies that softer monetary policy is relatively less expansionary for treated households in the downturn.

It is useful to note that the dependent variable in this subsection differs from the one used in our main macroprudential analysis. In our view, this is natural. The tightening of the LTV cap is

a one-time policy shock that permanently shifts household balance sheets and therefore the level of consumption over time. Monetary policy shocks, by contrast, arrive repeatedly and often offset one another, so their effects are more appropriately studied over shorter horizons. For this reason, consumption levels are more natural for the macroprudential analysis, while a forward-looking one-year consumption response is more natural for the monetary policy analysis. We confirm in the appendix, however, that the results form the previous section are robust to this alternative dependent variable.[7]

7.3 Results

Table 7 reports the main results. Columns 1 and 2 use the narrative monetary policy shock from Holm et al. (2021). Columns 3 and 4 use high-frequency shocks estimated for Norway by Brubakk et al. (2022). Columns 5 and 6 use actual changes in the policy rate. In columns 2, 4, and 6 we absorb year-quarter fixed effects.

Across all specifications, the coefficient on the quadruple interaction $Post_q \times Rogaland_i \times T_i \times MP_q$ is positive and statistically significant, at least at the 5% level. This is the key result of the section. It implies that expansionary monetary policy has a relatively weaker effect on consumption for households that were exposed to the tighter LTV cap and are located in the recession region during the downturn. Put differently, lower leverage induced by tighter macroprudential policy attenuates the transmission of monetary policy to household spending in bad times.

The point estimates are economically meaningful. We scale both the dependent variable and the monetary policy shocks by their overall standard deviations, so that the coefficients can be interpreted as the response of consumption (in standard deviations) to a one standard deviation monetary policy shock. In columns 1 and 2, which use the narrative shock, the coefficient on the quadruple interaction is around 0.071 and significant at the 5% level. In columns 3 and 4, which use the high-frequency shock, the estimate is 0.057 and again significant at the 5% level. In columns 5 and 6, which use actual policy rate changes, the estimate is about 0.035 and significant at the 1% level. The consistency of the sign and significance across three alternative shock measures suggests that the result is not driven by the choice of monetary policy series.

The coefficient on MP is negative and statistically significant in the specifications without year-quarter fixed effects, indicating that softer monetary policy generally raises consumption. This is exactly what one would expect given the local projections evidence above. Once year-quarter fixed effects are included, the main effect of the aggregate shock is absorbed, and the identifying variation comes from its interaction with treatment and region indicators.

[7] Specifically, we show that our main macroprudential results remain robust when we use the analogous forward-looking dependent variable at annual frequency, $(C_{i,t+1} - C_{i,t-1})/\overline{C}_{t-1}$, instead of the baseline normalized consumption measure: both the dynamic recession-effect (Figure A.3) and treatment-effect (Figure A.4) plots remain robust under this alternative outcome. Because the dependent variable is forward-looking, the event-time coefficients are shifted by one year relative to the corresponding Figures 3 and 4.

	(1) $\widetilde{C}$	(2) $\widetilde{C}$	(3) $\widetilde{C}$	(4) $\widetilde{C}$	(5) $\widetilde{C}$	(6) $\widetilde{C}$
Post × Rogaland	-0.076**	-0.076**	-0.088***	-0.087**	-0.086***	-0.085***
	(0.027)	(0.027)	(0.031)	(0.031)	(0.028)	(0.028)
Post × T	-0.026	-0.026	0.015	0.015	0.010	0.010
	(0.027)	(0.027)	(0.019)	(0.019)	(0.022)	(0.022)
Rogaland × T	0.182	0.181	0.150	0.148	0.164	0.160
	(0.135)	(0.134)	(0.138)	(0.137)	(0.137)	(0.136)
Post × Rogaland × T	0.061	0.061	0.095**	0.095**	0.077*	0.077*
	(0.038)	(0.038)	(0.042)	(0.042)	(0.040)	(0.040)
MP	-0.045***		-0.020**		-0.011*	
	(0.008)		(0.009)		(0.006)	
T × MP	-0.056*	-0.056*	-0.085***	-0.085***	-0.088***	-0.088***
	(0.030)	(0.030)	(0.021)	(0.021)	(0.007)	(0.007)
Post × T × MP	0.049	0.049	0.083***	0.083***	0.091***	0.091***
	(0.030)	(0.030)	(0.021)	(0.021)	(0.010)	(0.010)
Post × Rogaland × T × MP	0.071**	0.071**	0.057**	0.057**	0.035***	0.035***
	(0.028)	(0.028)	(0.025)	(0.025)	(0.010)	(0.010)
HH FE	Yes	Yes	Yes	Yes	Yes	Yes
YQ FE	No	Yes	No	Yes	No	Yes
Balance sheet controls	Yes	Yes	Yes	Yes	Yes	Yes
Sectoral income controls	Yes	Yes	Yes	Yes	Yes	Yes
Shock	Narrative	Narrative	HF	HF	Policy r	Policy r
R2	0.132	0.134	0.132	0.134	0.132	0.134
N	402,980	402,980	402,980	402,980	402,980	402,980

Table 7: Monetary policy transmission and LTV regulation

Notes: The dependent variable is the forward-looking one-year consumption measure defined in equation (11). $Post_q = 1$ in the recession period. $T_i = 1$ for treated households, i.e. potential house buyers in 2012. $T_i = 0$ for control households, i.e. potential house buyers in 2011. Narrative shocks are from Holm et al. (2021). High-frequency (HF) shocks are from Brubakk et al. (2022). Policy r denotes the actual quarterly change in the policy rate. All specifications include the same balance sheet and sectoral income controls as in the main consumption regressions, interacted flexibly with the post, region, and shock dimensions.

The coefficient on $Post_q \times Rogaland_i \times T_i$ is positive in all columns and statistically significant or close to conventional significance levels, in line with our baseline findings for macroprudential policy. Thus, even with this alternative, forward-looking dependent variable, treated households in Rogaland display smaller consumption declines during the recession. The recession effect however, i.e. the coefficient on $Post_q \times Rogaland_i$ is negative but not statistically significant. This is because the share of income coming from the oil sector, interacted with the post-dummy, is included in the regression. This absorbs much of the recession effect, as the average oil income share in Rogaland

is higher than in the rest of the country.

Taken together, the results suggest that tighter macroprudential policy makes households more resilient during the downturn, which in turn weakens the transmission of monetary policy. This is a resilience–effectiveness tradeoff: ex-ante macroprudential tightening dampens the contraction in consumption during downturns, making monetary stimulus both less necessary and less powerful.

8 Conclusion

This paper provides evidence that macroprudential regulation stabilizes household consumption during economic downturns, and documents a resilience–effectiveness tradeoff with monetary policy. We exploit Norway's tightening of the LTV cap from 90% to 85% in 2012 and the subsequent regional recession caused by the 2014 oil price collapse, comparing prospective first-time home buyers immediately before and after the policy change.

We document three main results. First, the LTV tightening reduces household leverage primarily through the extensive margin. The probability of purchasing a home falls by 4.0 percentage points (28% relative to the control group), translating into a 21% reduction in debt accumulation among prospective buyers. Among households that do purchase, the LTV distribution shifts visibly at the cap, but the aggregate effect on leverage conditional on buying is small. The extensive margin dominates the balance sheet response.

Second, these policy-induced leverage reductions translate into substantially smaller consumption declines during the recession. Consumption in the affected region falls approximately 8% relative to the rest of Norway, but the decline is only about 3% for households subject to the stricter LTV cap — a 60% attenuation. The effect is robust to household fixed effects, municipality-by-year fixed effects, and controls for balance sheet positions and sectoral income exposure, and we confirm parallel pre-trends. The consumption stabilization concentrates in durable and housing-related categories — transportation, utilities, and recreation — where the recession hits hardest, while non-durable necessities like food and services show minimal effects.

Using a principal stratification framework, we further show that the extensive margin accounts for essentially the entire consumption treatment effect. Under maintained assumptions of monotonicity, no effect on always-renters, and a resilience ordering on baseline consumption, the intensive-margin contribution is bounded to at most six percent of the total effect. LTV caps thus stabilize consumption during downturns primarily by keeping marginal households from becoming homeowners, rather than by improving the balance sheets of those who purchase.

Third, we document a resilience–effectiveness tradeoff in the interaction between macroprudential and monetary policy. Expansionary monetary policy significantly increases consumption during the recession, but the effect is substantially attenuated for households exposed to the tighter LTV cap. This result is consistent across narrative shocks, high-frequency shocks, and actual policy rate changes. At the same time, treated households require less monetary stimulus because their initial

consumption decline is smaller. Macroprudential policy thus reduces both the need for and the effectiveness of monetary stimulus during recessions.

Our findings have implications for the design of macroprudential policy. LTV caps stabilize household demand during downturns, but at the cost of weakening monetary policy transmission — fewer households carry mortgage debt that responds to interest rate changes. Whether this tradeoff favors tighter ex-ante regulation depends on the constraints facing monetary policy. Near the effective lower bound, the enhanced household resilience is particularly valuable because it reduces both the severity of initial shocks and the required monetary response. When monetary policy has ample room to adjust, the reduced transmission represents a clearer cost. Our evidence suggests that the probability of hitting the lower bound itself depends on the stance of macroprudential policy, as tighter regulation dampens the consumption declines that create the need for aggressive monetary easing.

Several limitations of our analysis point to directions for future work. Our setting involves a single policy change in one country, and the recession is regional rather than economy-wide, which limits the scope for general equilibrium effects and may affect external validity. Finally, while our decomposition shows that the extensive margin dominates, a fuller welfare analysis would require modeling the costs of homeownership foregone — including foregone housing consumption and wealth accumulation — against the benefits of consumption stability during downturns. We leave these questions for future research.

References

Aastveit, K. A., T. M. Fastbø, E. Granziera, K. S. Paulsen, and K. N. Torstensen (2024). Nowcasting norwegian household consumption with debit card transaction data. *Journal of Applied Econometrics 39*(7), 1220–1244.

Aastveit, K. A., R. E. Juelsrud, and E. G. Wold (2024). Does mortgage regulation stabilize household consumption? *Working Paper*.

Acharya, V. V., K. Bergant, M. Crosignani, T. Eisert, and F. J. McCann (2020). The real effects of the bank lending channel. *Journal of Monetary Economics 115*, 162–179.

Alpanda, S. and S. Zubairy (2019). Household debt overhang and transmission of monetary policy. *Journal of Money, Credit and Banking 51*(5), 1265–1307.

Arena, M., T. Chen, S. M. Choi, N. Geng, C. A. Gueye, T. Lybek, E. Papageorgiou, and Y. S. Zhang (2020). *Macroprudential policies and house prices in Europe*. International Monetary Fund.

Bank of Canada (2024). Macroprudential policy and monetary policy: Interactions and conflicts? Technical report, Bank of Canada.

Bergant, K., S. Gatti, A. Klemm, and R. Sahay (2023). Taming the tides of capital: How does macroprudential policy affect capital flows? *IMF Working Paper*.

Bernanke, B. S. (2011). Implementing a macroprudential approach to supervision and regulation. In *Federal Reserve Bank of Chicago Proceedings*, Number 1110.

Brubakk, L., S. ter Ellen, Ørjan Robstad, and H. Xu (2022). The macroeconomic effects of forward communication. *Journal of International Money and Finance 120*, 102536.

Carroll, C. D., M. Otsuka, and J. Slacalek (2011). The dynamics of consumer spending: The role of durable goods, information, and liquidity constraints. *Economics Letters 111*(1), 99–102.

Cerutti, E., S. Claessens, and L. Laeven (2017). Changes in prudential policy instruments—a new cross-country database. *International Journal of Central Banking 13*(1), 477–503.

Cloyne, J., C. Ferreira, and P. Surico (2020). Monetary policy when households have debt: new evidence on the transmission mechanism. *The Review of Economic Studies 87*(1), 102–129.

DeFusco, A. A., S. Johnson, and J. Mondragon (2020). Regulating household leverage. *The Review of Economic Studies 87*(2), 914–958.

Farhi, E. and I. Werning (2016). A theory of macroprudential policies in the presence of nominal rigidities. *Econometrica 84*(5), 1645–1704.

Flodén, M., M. Kilström, J. Sigurdsson, and R. Vestman (2021). Household debt and monetary policy: Revealing the cash-flow channel. *The Economic Journal 131*(636), 1742–1771.

Galaasen, S., M. Mæhlum, and S. Ahn (2024). The cash-flow channel of monetary policy: Evidence from billions of transactions. *Available at SSRN 5089327*.

Greenwald, D. L. (2018). The mortgage credit channel of macroeconomic transmission. *MIT Sloan Research Paper* (5184-16).

Harding, M. and M. Klein (2022). Resolving the missing deflation puzzle. *Journal of Monetary Economics 126*, 15–34.

Holm, M. B., P. Paul, and A. Tischbirek (2021). The transmission of monetary policy under the microscope. *Journal of Political Economy 129*(10), 2861–2904.

International Monetary Fund (2013). Key aspects of macroprudential policy. Technical report, International Monetary Fund.

International Monetary Fund (2017, April). Global financial stability report: Getting the policy mix right. Technical report, International Monetary Fund.

International Monetary Fund (2024, April). World economic outlook: Steady but slow: Resilience amid divergence. Technical report, International Monetary Fund.

Jiménez, G., S. Ongena, J.-L. Peydró, and J. Saurina (2017). Macroprudential policy, countercyclical bank capital buffers, and credit supply: evidence from the spanish dynamic provisioning experiments. *Journal of Political Economy 125*(6), 2126–2177.

Jones, C., V. Midrigan, and T. Philippon (2022). Household leverage and the recession. *Econometrica 90*(5), 2471–2505.

Juelsrud, R. E. and E. G. Wold (2025). The importance of unemployment risk for individual savings. *Management Science 71*(8), 6746–6769.

Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary policy according to hank. *American Economic Review 108*(3), 697–743.

Kuttner, K. N. and I. Shim (2016). Understanding the effects of macroprudential policy: An empirical investigation. *Journal of Money, Credit and Banking 48*(5), 769–796.

Martín, A., C. Mendicino, and A. Van der Ghote (2021). The coordination of prudential and monetary policies. *Working Paper*.

Mian, A., K. Rao, and A. Sufi (2013). Household balance sheets, consumption, and the economic slump. *The Quarterly Journal of Economics 128*(4), 1687–1726.

Mian, A. and A. Sufi (2011). House prices, home equity-based borrowing, and the us household leverage crisis. *American Economic Review 101*(5), 2132–2156.

Mian, A. R., A. Sufi, and E. Verner (2017). Household debt and business cycles worldwide. *The Quarterly Journal of Economics 132*(4), 1755–1817.

van Bekkum, S., M. Gabarro, R. M. Irani, and J.-L. Peydró (2024). The real effects of borrower-based macroprudential policy: Evidence from administrative household-level data. *Journal of Financial Economics 152*, 103763.

Van der Ghote, A. (2021). Coordinating monetary and financial regulatory policies. *American Economic Journal: Macroeconomics 13*(2), 20–57.

A Additional figures

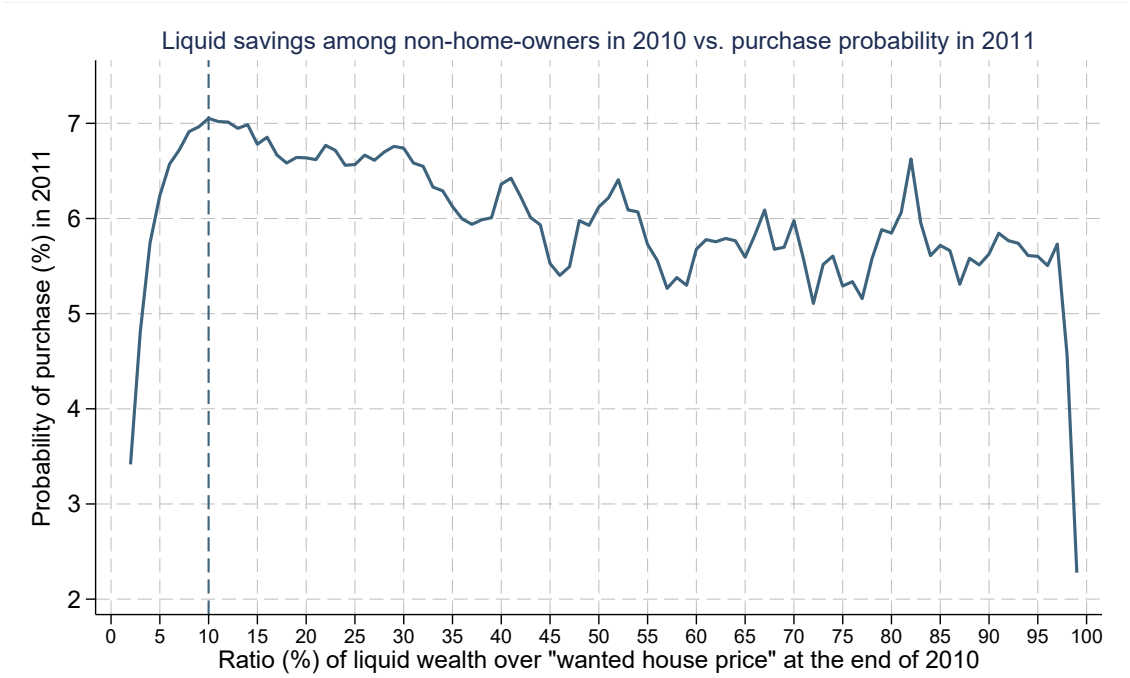


Figure A.1: Liquid wealth and probability of housing purchase

Notes: The figure shows the relationship between non-homeowners' liquid wealth relative to their "wanted house price" (as defined in Section 4 strategy) at the end of 2010 and their average probability of purchasing a home in 2011. The probability is computed as the number of purchasing households divided by the total number of households (purchasing and non-purchasing). The plotted line is smoothed using a moving window of 5 percentage points (i.e., ±2 percentage points on each side) around each value on the x-axis.

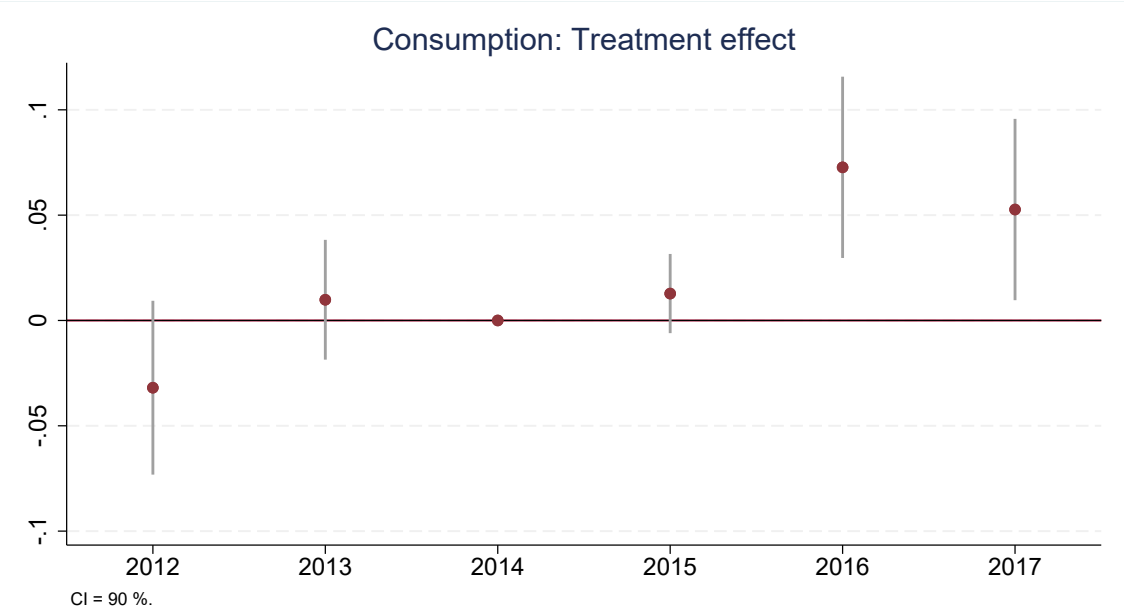


Figure A.2: Treatment effect: Estimated Γ_3-coefficients from equation (1)

Notes: Consumption dynamics in Rogaland for treated households. Consumption is defined as $\overline{C}_{it} = C_{i,t}/C_{i,2014}$. $T_i = 1$ for treated households, i.e. potential house buyers in 2012. $T_i = 0$ for control households, i.e. potential house buyers in 2011. No control variables.

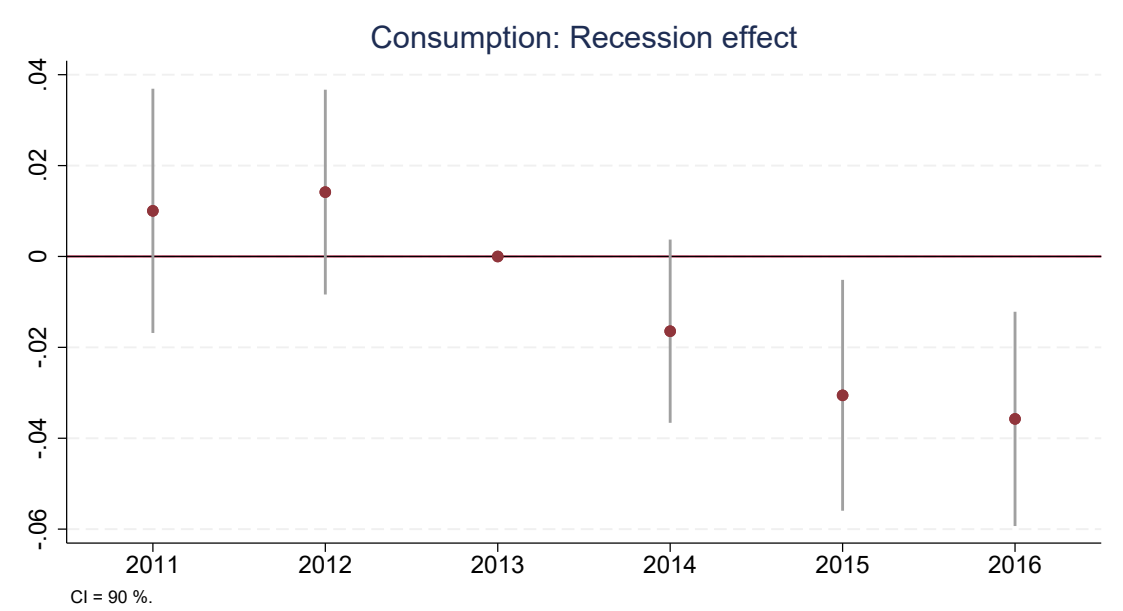


Figure A.3: Recession effect with forward-looking consumption measure: Estimated Γ_1-coefficients from equation (1)

Notes: Consumption dynamics in Rogaland using the forward-looking annual consumption measure. The dependent variable is defined as $\widetilde{C}_{it} = (C_{i,t+1} - C_{i,t-1})/\overline{C}_{t-1}$, where $\overline{C}_{t-1}$ is average household consumption in year $t-1$. $T_i = 1$ for treated households, i.e. potential house buyers in 2012. $T_i = 0$ for control households, i.e. potential house buyers in 2011. Because the dependent variable is forward-looking, the event-time coefficients are shifted by one year relative to Figure 3.

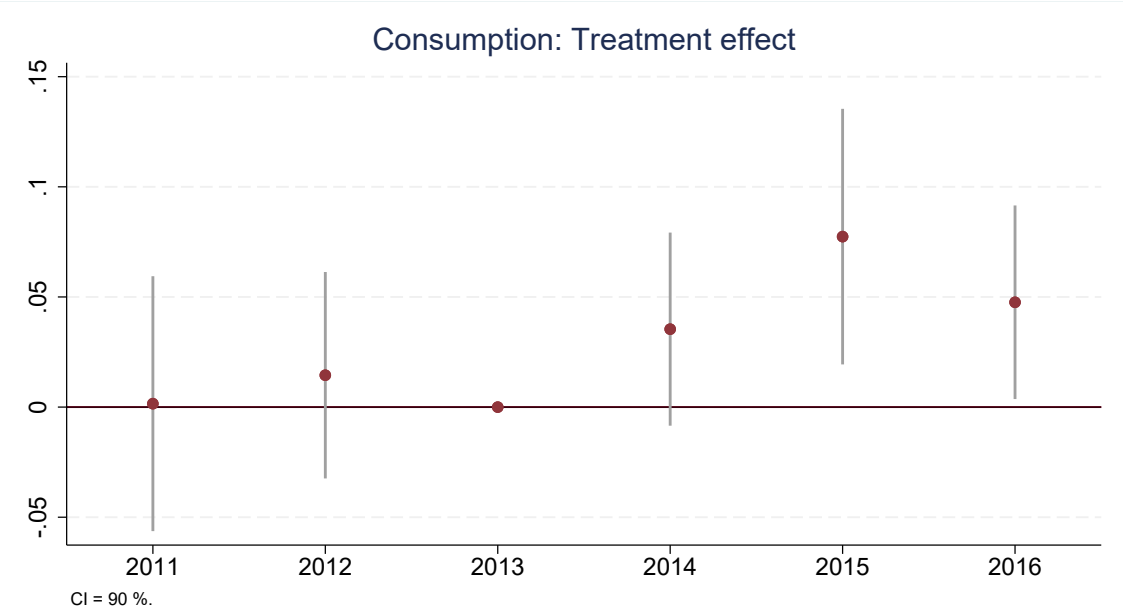


Figure A.4: Treatment effect with forward-looking consumption measure: Estimated Γ_3-coefficients from equation (1)

Notes: Consumption dynamics in Rogaland for treated households using the forward-looking annual consumption measure. The dependent variable is defined as $\widetilde{C}_{it} = (C_{i,t+1} - C_{i,t-1})/\overline{C}_{t-1}$, where $\overline{C}_{t-1}$ is average household consumption in year $t-1$. $T_i = 1$ for treated households, i.e. potential house buyers in 2012. $T_i = 0$ for control households, i.e. potential house buyers in 2011. Control variables include 2014-values of age, income, wealth, parental wealth, and sectoral income shares, interacted with year and Rogaland-dummy. Because the dependent variable is forward-looking, the event-time coefficients are shifted by one year relative to Figure 4.

B Additional tables

	(1) Purchase price	(2) Financial wealth	(3) Financial wealth	(4) Financial wealth
T	-109244***	68754**	66719**	-97377
	(40384)	(33143)	(33187)	(85280)
Buy			-60330	
			(50721)	
Constant	2211437***	400218***	408585***	418605***
	(26041)	(23387)	(24422)	(55838)
Sample	Buyers	All	All	Buyers
Observations	4,122	32,801	32,801	3,998

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table B.1: Balance sheet effects of LTV-tightening: purchase price and financial wealth

Notes: House purchase price and financial wealth – in year of potential buyer. $T_i = 1$ for treated households, i.e. potential house buyers in 2012. $T_i = 0$ for control households, i.e. potential house buyers in 2011.