Household Debt, Consumption and Inequality

Berrak Bahadir
Department of Economics
University of North Carolina at Greensboro
b_bahadi@uncg.edu

Kuhelika De
Department of Economics
Grand Valley State University
dek@gvsu.edu

William D. Lastrapes
Department of Economics
University of Georgia
last@uga.edu

September 3, 2019

Abstract

This papers examines the potential link between household credit shocks and income inequality at the national level. For a sample of 32 developed and developing countries, we show that aggregate consumption temporarily increases in the short run and decreases in the long run in the face of credit shocks, and that this dynamic response is larger in absolute value, in the short and long run, for countries with high income inequality compared to those with low income equality. We develop a simple dynamic theoretical model, based on binding credit constraints, to illustrate the model's mechanisms.

Keywords: Credit constraints, credit shocks, income distribution, VAR, Gini coefficient

JEL codes: E21, E32, E44, E51

1 Introduction

The recent work of Mian, Sufi, and Verner (2017) and others documents the importance of household debt as a driver of business cycles for developed and emerging-market economies. Household debt appears to work through a 'household demand channel' (Mian, Sufi, and Verner, 2019) and affects both the boom and bust of a debt cycle. During the boom, household borrowing increases consumption and contributes to an increase in economic activity; but such borrowing ultimately brings about a bust as households retrench in the face of mounting debt. A consensus has emerged that household debt can generate short-term gains but at a cost of significant reductions in medium- to long-term growth. We do not yet, however, have a complete picture of the economic mechanisms at play.

Our paper aims to add to this picture by looking at the relationship between household debt and aggregate consumption from a different angle than most previous studies. In particular, we examine the extent to which income inequality contributes to the household demand channel in response to household credit shocks. Section 2 of the paper takes this question to the aggregate data. We first estimate the dynamic effects of household credit shocks on aggregate consumption for a sample of 32 countries, using standard VAR techniques and treating countries individually. We show that household credit shocks tend to have positive effects on consumption – which most likely indicates the importance of binding credit constraints at the aggregate level – but that these effects die out over time and for some countries eventually become negative. This finding is consistent with most of the related literature. We then run cross-country regressions to gauge the effect of income inequality, as measured by country-specific Gini coefficients, on the sensitivity of consumption to household credit shocks estimated in the first stage of the data analysis. We show that countries having higher Gini coefficients, and thus more unequal income distributions, than other countries exhibit greater short-run gain and greater medium- to long-run pain from household credit shocks. Thus, we find an empirical link between income inequality and the household demand channel.

Section 3 of the paper develops a simple theoretical model to illustrate how income inequality can generate these findings. The model relies on financial market imperfections and credit constraints to motivate a link between inequality and the incidence of household credit shocks. In the model, the burden of credit constraints depends on a household's income – high-income earners never face a borrowing constraint and low-income earners always do. Middle-earners, however, may be credit constrained (like low-earners) if their income level is sufficiently low, but they can be unconstrained (like high-earners) if their income is sufficiently high. The model implies that if the income share of middle-earners falls, the country's income distribution becomes more unequal (the Gini coefficient rises) at the same time that there is an increase in the number of households that are credit constrained. This increase in the incidence of credit constraints further implies that credit shocks have larger aggregate effects on consumption the greater is income inequality. We simulate the model to show that this mechanism can explain the shapes and magnitudes of the impulse response functions estimated from the data.

Our work contributes to the vast literature on understanding the effects of credit supply shocks on the overall economy. In addition to the two papers cited in the first paragraph, an abbreviated list of related research includes Mian and Sufi (2018), Justinian, Primiceri, and Tambalotti (2015), Favilukis, Ludvigson, and Van Nieuwerburgh (2017), Gennaioli, Shleifer, and Vishny (2012), Schmitt-Grohé and Uribe (2016), Korinek and Simsek (2016), Farhi and Werning (2016), Bahadir and Gumus (2016), Cloyne et al. (2019) and Abdallah and Lastrapes (2012). Only a handful of papers in this area specifically consider income distribution and inequality as an important factor for the effects of debt. In a strictly empirical study, Alter, Feng, and Valckx (2018) show that the mortgage-debt share of lower income households – as a measure of unequal access to financial markets – affects the relationship between household debt and growth. However, that paper neither examines other measures of inequality nor links the findings to specific theory. Kumhof, Ranciere, and Winant (2015) show that higher leverage and crises arise endogenously in response to a growing share of high-income households, but do not account for the important difference – documented in much of the studies noted above – between household and firm debt. Iacoviello (2008) shows that the prolonged rise in household debt in the US can be explained only by the concurrent increase in income inequality, a finding supported by Kumhof, Ranciere, and Winant (2015), although our paper is agnostic about whether economies with high income inequality experience faster growth in household debt.

2 Data analysis

2.1 Estimating the dynamic response of consumption to credit shocks

In the first part of the data analysis we estimate country-specific dynamic responses of real aggregate consumption to household credit shocks. We do so in the baseline case by inverting estimated VAR models to obtain impulse response functions, and then checking for robustness to model misspecification using the local projections approach of Jorda (2005) to estimate impulse response functions directly. In the baseline model we separately estimate, for each of the 32 countries in our sample, a VAR model that includes the log of real consumption, c_t, the ratio of non-financial firm debt to GDP, $D_t^f = \frac{F_t}{Y_t}$, and the ratio of household debt to GDP, $D_t^h = \frac{H_t}{Y_t}$. We use the same variables in the local projections estimations. We choose this three-variable system to be comparable to Mian, Sufi, and Verner (2017).[1] The data are seasonally-adjusted quarterly observations over a sample period from 1990 to 2017, where the first and last observations vary within this range across countries based on data availability.[2] Table 1 lists the countries in the sample, time sample ranges, the log of real consumption, the debt ratios, and two other variables – a financial development

[1]Estimating separate VAR models for each country in the sample is less restrictive than the panel approach of Mian, Sufi, and Verner (2017). In effect, our approach is tantamount to estimating a panel model that includes fixed effects dummy variables that interact with ALL lagged right-hand-side variables in the system.

[2]Household and non-financial firm debt are from the "Long series on credit to the private non-financial sector" database of the Bank for International Settlement. GDP, household consumption and Consumer Price Index (CPI) series are from the International Financial Statistics (IFS) Database. We use quarterly data in current and constant prices from the IFS and deflate nominal consumption spending by the CPI to get real consumption.

index and the gini coefficient – that will be used in the cross-sectional regressions below.

Because we choose to focus on the impulse responses to a household credit shock only (and not firm-credit shocks or consumption shocks), we need not fully identify the empirical model. Instead we impose only two identifying restrictions that are sufficient to just-identify the structural shocks of interest: we assume that neither consumption nor the firm debt ratio responds contemporaneously to household credit shocks. The assumption that consumption is slow to respond to credit shocks is standard in the related literature, see for example Mian, Sufi, and Verner (2017, p. 1764). The second restriction defines a *household* credit shock as one that has an immediate effect on household debt – either the demand to borrow by or the supply of funds to households, such as a loosening or tightening of borrowing limits for consumer credit – but that does not directly affect, on impact, the demand for or supply of *firm* debt. These restrictions can be implemented using the standard Cholesky decomposition of the reduced form residual covariance matrix estimated from the VAR, using the ordering noted above. Thus, our results are based on a minimal set of identifying restrictions and should therefore be consistent with a wide range of theoretical models.[3] Although our approach to identification is standard, the appendix provides a brief explanation of the restrictions for both the VAR and local projection models.

We do not attempt at this stage to precisely identify whether household credit shocks in our model and data are demand or supply induced. Interest rate variation provides especially useful identifying information in this context; however, we do not have sufficient data on interest rates to perform a convincing analysis. We rely on the findings of Mian, Sufi, and Verner (2017) that credit supply shocks, as opposed to demand, are most likely

[3]As long as the household debt variable is ordered last, identification of the responses of each variable to household credit shocks does not depend on the ordering of the first two variables. Note that to identify firm credit shocks using an analogous strategy would require re-ordering firm credit as the last variable in the Cholesky ordering, so that such a shock would have no contemporaneous effect on *household* credit. Simply interpreting the shock from the second equation from the original ordering as a firm credit shock is misleading since its impact effects on both firm and household debt would not be restricted in this case. See Mumtaz, Pinter, and Theodoridis (2018) for a systematic attempt to assess credit shock identification in a structural VAR framework. Recursively-identified VAR models do not fare well in their paper, but they do not separate household from firm credit, so their findings are not conclusive evidence against our approach.

driving our results.

We report the impulse response functions estimated from our baseline VAR models for each country in Figures 1 and 2, which plot the dynamic responses of both the log of real consumption (c_t) and the household credit ratio (D_t^h) to a household credit shock up to a 24-quarter horizon.[4] The first figure shows responses to a unit shock, the second to a standard deviation shock. We look at the first case to compare responses across countries to shocks of a common magnitude. The second case accounts for potential differences in the scale of credit shocks across countries. Each VAR includes four lags of the system variables and a deterministic time trend, which is sufficient to whiten the residuals.

By construction the immediate effect of a unit household credit shock, as seen in the first figure, is to increase the household credit ratio by one percentage-point on impact for each country (red curve), but the data determine the estimated dynamics of the response over the remaining horizons. For most countries the response of the household credit ratio is persistently positive over the short- to medium-run, but the degree of persistence varies widely across countries. For countries like Australia, Germany and the US, household credit as a fraction of GDP remains well above its initial steady-state up to 24 quarters after the shock; for others, like Japan, Russia and Switzerland, household credit relative to GDP rises in the short-run but falls below its initial steady state after two to three years.

There is variation in the consumption response across countries as well (blue curve). Our recursive identification scheme forces the impact effect of consumption to be zero for all countries, but the data show that consumption tends to rise in the short-run beyond the impact horizon. Indeed, for 24 of the 32 countries in the sample the consumption response is positive at some point over the first two years after the shock. For the US, the maximum consumption response to a household credit shock that initially increases household credit by 10 percentage points (for example, from the US mean household credit/gdp of 76.8% to 86.8%) is 6%, which happens at the eight-quarter horizon. For the UK, the maximum

[4] Because we attempt to explain cross-country variation in the next stage, and to avoid graphical clutter, we do not report standard error bands in the figure.

consumption response to a local household credit shock of the same magnitude is around 8% for a similar horizon.

An evident pattern from the figure is that domestic household credit-to-GDP expansion in the emerging market economies in our sample – Argentina, Brazil, Mexico, Russia and Turkey – leads to a relatively large household spending response in the short run. Consumption in Argentina, for example, rises by almost 7% in response to a *one* percentage point household credit ratio shock, an order of magnitude more than the US consumption increase. The maximum responses of Brazil, Mexico, Russia and Turkey are in the range of 1.5 to 2.5%, all within the first year after the shock. Argentina and Russia also exhibit large busts in consumption over the medium run. In Russia, the credit boom leads to an ultimate decline in output of nearly 10%, which in Argentina is over 3% between three and four years after the shock. Consumption booms in the short run and busts in the long run in Italy and Greece are pronounced as well. Thailand experiences a large bust but a smaller boom than the other emerging market economies.

We investigate the source of this cross-country variation in more detail below, but one potential explanation for these differences is that a given percentage-point increase in the household debt-to-GDP ratio is relatively large for the developing nations given their small average household debt ratios. In Argentina, for example, this ratio is 5% on average over our sample, compared to 77% in the US (see Table 1). Average debt ratios in the other four countries range from 7% to 16%. A 10 percentage point increase in a country with a 10% debt-income ratio is a doubling of that ratio; the same percentage point increase for a country with an 80% ratio is only a 12.5% increase. This explanation is less plausible for Italy, Greece and Thailand, since their debt ratios range on average from 30% to 50%.

Figure 2 accounts for this difference in the scale of household credit across countries by normalizing on standard deviation shocks. By construction, the shapes of the response functions will be identical across the two figures, but the magnitudes measured along the vertical axis can differ. The figure shows that accounting for the estimated scale of credit shocks across countries does not alter our inference. For example, the consumption responses in Argentina and Brazil remain more than double the size of the US response, even though average shock size is smaller in the former countries.

Figure 3 summarizes the dynamic responses of consumption to the unit household credit shocks from the first figure to better illustrate the cross-country variation in those responses. The top panel plots the responses at horizons 2 (the period after the initial shock), 4, 8, 12, 16, 20 and 24 quarters. The solid red curve is the cross-sectional mean response across the sample of 32 countries for each horizon from 1 to 24. The mean positive response in the short run and negative response over the medium to long run are consistent with the boom-bust hypothesis of household credit shocks, and are similar to the aggregate responses of GDP as reported by Mian, Sufi, and Verner (2017, Figure 1, p. 1765). The extreme values in Argentina and Russia are evident in the graph, but there is substantial variation across the responses of the other countries as well. The bottom panel contains each country's peak response, at the horizon at which that peak occurs (with the cross-sectional mean again superimposed).

Figures 4 and 5 show the same set of results for the impulse response functions estimated directly using local projections. As is typical, the local projections response functions are not as smooth as those from the VAR; however, shapes and magnitudes are generally comparable across the two methods.[5]

2.2 Estimating cross-country variation in the response of consumption

In this sub-section we attempt to explain the estimated cross-country variation in boom-bust dynamics of consumption in response to household credit shocks from the previous analysis, focusing on the potential role of income inequality. We estimate the cross-sectional regression model

$$y_i = \beta_0 + \beta_1 z_{i1} + \beta_2 z_{i2} + \epsilon_i \tag{1}$$

[5]See Barnichon and Brownlees (2019) for an approach to smooth local projection impulse responses.

for countries $i = 1, \cdots, 32$, where the dependent variable y_i is a summary measure of each country's dynamic response of consumption to household credit shocks estimated above. The explanatory variables z_{i1} and z_{i2} are proxies for nation-wide financial development, included as a fundamental control variable, and income inequality, the variable of primary interest. Although the impulse responses for the dependent variable are generated from first stage VAR estimation, measurement and specification error in that stage will be captured by the cross-sectional regression error term and will affect inference only to the extent that measurement error is correlated with the primary explanatory variables. There are no obvious reasons to expect such correlation. β_1 and β_2 measure the marginal effects of financial development and inequality on the estimated dynamic responses; our primary focus is on β_2.

Our two-stage approach to explaining cross-sectional variation in dynamic responses is more general than conventional panel data methods, since we do not impose the potentially severe constraint that parameters are identical across countries. The approach has precedents in the literature; see for example Lastrapes and McMillin (2004), Cecchetti (1999), and Aizenman et al. (2019), among many others.

Because we are interested in the short-run boom in consumption and the medium- to long-run bust, we run the cross-country regression separately for various measures of the dependent variable, y_i. In particular, we alternately set y_i to be the impulse response coefficient estimates on 'impact' (with a one-period lag) of consumption to the credit shock ($c_{qj}, j = 2$) and for each four-quarter horizon up to quarter 24. We also consider the maximum response over the first twelve quarters (c_{max}), and cumulative responses over the short-run (1 to 12 quarters) and medium- to long-run (12 to 24 quarters). Table 2 reports for each country the maximum response over the first twelve quarters (c_{max}), and responses at quarters 2, 4, 8 ($c_{qj}, j = 2, 4, 8$).

To measure the level of a country's financial development (z_1) we use the index developed by Svirydzenka (2016), which combines information on the depth, access and efficiency of financial institutions and markets in that country. A higher value for z_1 indicates a higher level of financial development. For inequality (z_2) we use a country's Gini index, obtained from the Standardized World Income Inequality Database (Solt, 2016). A higher value for z_2 indicates a higher level of income inequality. For each series we use the average value for each country over its sample range. Table 1 reports these average values.

Our prior view is that the extent to which households face binding credit constraints plays an important role in driving both the boom and bust after a consumer credit shock, and that inequality and the extent to which credit constraints bind are linked. We develop a more formal model of this mechanism in the following section, but here generally describe the implications for our cross-sectional regression. We would expect countries with a high degree of income inequality – and therefore high Gini coefficients – to have a larger number of households in the low-income tails of the distribution and thus to be credit constrained than those countries with less income inequality. Thus, credit supply shocks are more likely to affect the extent of binding credit constraints in more unequal economies and therefore to elicit a larger consumption response. Likewise, we would expect lower financial development to be associated with a greater extent of binding credit constraints; thus, in financially undeveloped economies household spending will exhibit a relatively large response to a credit shock compared to countries with low financial development and a high degree of binding credit constraints.

This argument implies that consumption booms and busts due to credit shocks will be large in countries with relatively low financial development and high income inequality. Consider a positive shock to domestic household credit supply in two countries that differ only according to income inequality and financial development. In the short run, consumption in both economies expands in response to the shock because credit constraints bind at least for some households; however, all else the same it will rise more in the low development/high inequality economy because the credit loosening affects a larger number of individuals. In terms of our cross-sectional regression model, a negative value for β_1 and a positive value for β_2 for short-run horizons would be consistent with this story. As the effects of the positive shock unwind, perhaps because of excess lending and other mechanisms described by Mian, Sufi, and Verner (2017), the expansionary phase makes way for the contractionary phase and consumption falls. We would again expect this bust to be larger for the low development/high inequality countries. Because this contraction happens over the medium-to long-run, we would expect to see positive β_1 estimates and negative β_2 estimates over those longer horizons.

This pattern is precisely what we find in the data, as reported in Table 3 for the baseline VAR results with respect to both unit and standard deviation credit shocks. For the impact, four-quarter and short-run maximum effects, our estimate of β_1 is negative and of β_2 is positive, both statistically significant, for both shock scalings. The absolute magnitude of the effect is, however, larger for the unit shocks compared to standard deviation shocks. For the unit shocks in the period after impact, a standard deviation increase in the financial development index of 0.1525 reduces the semi-elasticity of consumption by 53 basis points, which is almost one-half of the the cross-country standard deviation of the semi-elasticity of 120 basis points. A standard deviation increase in equality leads to a short-run increase of 49 basis points in c_{q2}, also economically important. Note as well that the adjusted R^2 is over 0.40 in the early horizons, which means that these two variables alone explain almost half the cross-country variation in the impulse responses at the short-run horizons. For unit shocks these signs reverse at longer horizons: β_1 is significantly positive after eight quarters, while β_2 becomes negative but not statistically significant. When we look at the cumulative consumption responses over horizons 1 to 12, β_2 is 0.41, which is large and statistically significant, whereas this coefficient is estimated to be -0.25 (though not statistically different from zero) for horizons 12-24, which is consistent with our priors. The final two rows of the table set y_i to be the cumulative response over a short-run period less the cumulative response over the medium- to long-run horizon. The idea here is to measure variation in the extent of the boom and bust in terms of how far the response falls from the short-run to the long-run (the amplitude of the cycle). We again find a statistically significant negative effect of financial development (less developed countries have bigger bust in the consumption response) and a positive effect of inequality (more inequality leads to a larger bust).

These findings are roughly consistent with those found when using the local projection responses. They are also consistent with the findings of Alter, Feng, and Valckx (2018, Figure 7, p. 44) which show that countries in which credit participation for low income households is relatively high suffer less from negative credit shocks than countries with low participation rates.

Strict causal inference in this case requires that inequality be exogenous in the cross-sectional regression (conditional on the level of financial development). We find that our estimates are robust to some quick checks for omitted variable bias, and there is no obvious reason to suspect reverse causality. While joint dependence on unobservables might possibly bias the interpretation, our results present a *prima facie* case for causality, and at least document an important empirical association between inequality and the housing demand channel.

3 Theory

3.1 Overview

In this section we develop and calibrate a simple dynamic model to quantify the link between inequality and aggregate consumption's response to household debt shocks. Our modest objective is to quantitatively illustrate one potential mechanism that can explain this link; we do not attempt to construct a complete general equilibrium model, leaving that more ambitious goal for future research.

We work with a model of a small open-economy with incomplete financial markets. Our notation is for a single country, but we assume the same structure holds for all countries. There are three groups of infinitely-lived households: low-income earners with population share ω_l, middle-income earners with population share ω_m, and high-income earners with population share $\omega_h = 1 - \omega_l - \omega_m$. The economy's output is exogenous and given by an autoregressive stochastic process

$$y_t = (1 - \rho_y)\overline{y} + \rho_y y_{t-1} + \varepsilon_{y,t} \tag{2}$$

where $\epsilon_{y,t}$ is white noise and a bar above a variable denotes its steady-state value.[6] The shares of total income received by low-, middle-, and high-income earners are z_l, z_m, and $z_h = 1 - z_l - z_m$ respectively. Throughout the analysis we assume that population shares of the three groups remain constant over time. All domestic households are net borrowers; the source of domestic borrowing comes from lenders in international capital markets.

Credit constraints drive the link between inequality and spending sensitivity in our model. We assume that international lenders set an exogenous income threshold above which borrowers can borrow without any limits. Below this threshold, however, lenders impose quantity constraints based on expected income. The distinction between income groups lies in the expected burden of the credit constraints. We assume that low earners have anticipated income that never exceeds the threshold so they always face binding constraints, while high earners always have sufficient anticipated income to avoid binding constraints. On the other hand, we assume that middle-income earners have income levels that span the lending threshold. As this group's income level increases, relative to high-income earners, and average income rises above the lending threshold, a smaller share of the overall population faces binding credit constraints and consumption becomes less sensitive to credit shocks. At the same time, the rise in income share of middle earners lowers income inequality and thus the Gini coefficient.

[6] Kumhof, Ranciere, and Winant (2015) in a related context also assume that aggregate income is exogenous.

3.2 High-income earners

High-income earners maximize the expected lifetime utility function

$$E_0 \sum_{t=0}^{\infty} \beta_h^t \left(\frac{c_{h,t}^{1-\sigma}}{1-\sigma} \right), \tag{3}$$

where $\beta_h \in (0,1)$ is the discount factor, $c_{h,t}$ is consumption, and σ is the risk-aversion parameter for these households. This income group can borrow and lend without constraints, but faces a small convex financial intermediation or adjustment cost when borrowing at levels that are different from the steady-state. The budget constraint of high-income earners is thus

$$c_{h,t} + R_{t-1} b_{h,t-1} + \frac{\psi}{2}(b_{h,t} - \bar{b}_h)^2 = b_{h,t} + z_h y_t, \tag{4}$$

which holds for all periods in the planning horizon, and where $b_{h,t}$ denotes high-income household debt at time t and ψ is an adjustment cost parameter. R_{t-1} is the gross interest rate on debt that matures at time t and is taken to be exogenous and equal to the stochastic process for the world real interest rate. High-income earners maximize equation (3) with respect to (4), generating the optimality condition

$$\frac{c_{h,t+1}}{c_{h,t}} [1 - \psi(b_{h,t} - \bar{b}_h)] = \beta_h R_t \tag{5}$$

3.3 Low-income earners

Low-income earners' utility from consumption, $c_{l,t}$, takes the same functional form as high-income earners but they are more impatient and have a lower discount factor, $\beta_l < \beta_h$.[7]

[7] Impatience is a common assumption in the literature to obtain an equilibrium in which some agents are credit constrained (Iacoviello, 2005). Frederick, Loewenstein, and O'donoghue (2002) summarize the empirical evidence for discount rate heterogeneity across different types of households and Becker and Mulligan (1997) provide theoretical support for the hypothesis that the rich tend to be more patient.

They face the intertemporal budget constraint

$$c_{l,t} + R_{t-1} b_{l,t-1} = b_{l,t} + z_l y_t. \tag{6}$$

Low-income households also face a quantity constraint on their ability to borrow because their income is less than the lending threshold set in international capital markets; they are therefore less credit-worthy than top-tier income households. Their credit constraint is such that the total value of debt cannot exceed a time-varying fraction of expected income in the next period. As in Ludvigson (1999), we tie borrowing to expected future income because income is assumed by lenders to be associated with the borrower's financial health and ability to service the debt. The credit constraint of low-income earners takes the form

$$b_{l,t} \leq (1-\theta)\mu_t z_l E_t(y_{t+1}) + \theta b_{l,t-1}. \tag{7}$$

where $0 \leq \theta \leq 1$ measures the degree of inertia in the borrowing limit, which allows us to generate persistence in cycles. As $\theta \to 0$ the constraint takes the usual form and μ_t can be interpreted as the maximum loan-to-income ratio required by lenders. This specification is similar to the borrowing constraint used in Guerrieri and Iacoviello (2017), which includes a persistence term to reflect the slow adjustment of borrowing to house price changes. When calibrating the model we assume $\beta^l < 1/\bar{R}$, which guarantees that the credit constraint is binding in and around the steady state.

We assume μ_t follows the stochastic process

$$\mu_t = \bar{\mu} \exp(\tilde{\mu}_t) \tag{8}$$

$$\tilde{\mu}_t = \rho_\mu \tilde{\mu}_{t-1} + \varepsilon_{\mu,t}. \tag{9}$$

Low-income earners are households for whom $\bar{\mu} > 0$ always takes a finite value. A positive shock to $\epsilon_{\mu,t}$ can be interpreted as an unanticipated but persistent (depending on the value of the ρ_μ) loosening of the supply of credit. The variance of this shock is σ_μ^2.

3.4 Middle-income earners

Middle-income earners have identical preferences to high- or low-income earners. However, as noted above these earners have income that can either exceed or fall below the lending threshold set by international capital markets. We make this assumption operational by allowing the credit constraint for this sector to be state-dependent. In particular, we assume that the budget and credit constraints facing middle earners are

$$c_{mt} + R_{t-1}b_{m,t-1} + \frac{\psi(z_m)}{2}(b_{m,t} - \bar{b}_m)^2 = b_{m,t} + y_t z_m \tag{10}$$

$$b_{m,t} \leq (1-\theta)\mu_m(z_m)z_m E_t(y_{t+1}) + \theta b_{m,t-1} \tag{11}$$

$$\psi(z_m) = \begin{cases} \psi & z_m \geq \phi \\ 0 & z_m < \phi \end{cases} \tag{12}$$

$$\mu_m(z_m) = \begin{cases} \infty & z_m \geq \phi \\ \mu_t & z_m < \phi. \end{cases} \tag{13}$$

We take ϕ to be the income share that reflects the level of income consistent with the lending threshold. As income share z_m rises above ϕ, the middle-earner group becomes more like the high-income group and pays adjustment costs but faces no constraint (we can think of the 'loan-to-income' ratio as going to infinity in this case). At the same time, since the middle group's income share rises and its population share is fixed, the economy's Gini coefficient necessarily falls. For small z_m, there is no adjustment cost for borrowing, but middle earners now face the same borrowing constraint as low income households. Our simulation experiment below considers such a 'regime-shift' for middle earners.

To fix ideas about this experiment, consider a simple numerical example. Suppose that there are 1,000 households in the economy, with respective population, per-capita income and income shares:

i	n_i	y_i/n_i	z_i
l	20%	\$10,000	1.2%
m	70%	\$100,000	40.7%
h	10%	\$1,000,000	58.1%

Assuming per-capita income is the same for individuals within each group, the implied Gini coefficient is 55.5.[8] Now, increase income in the low-income sector by 20% and in the middle-income sector by 50%, holding income constant for the rich. The income share of middle earners rises to 50.6%, that of high earners falls to 48.2%, and the Gini coefficient falls to 47.5%. If the lending threshold for individual borrowers is \$125,000, then the increased level of income and income share of the middle-income group lead, in our model, to a reduction in the population share of credit-constrained households from 90% to 10%.

3.5 Equilibrium

In equilibrium, all households maximize their respective lifetime utilities with respect to the relevant credit and budget constraints, the market for borrowing and lending clears, and the market clearing condition for goods holds. The market-clearing conditions are

$$B_t = (1 - \omega_m - \omega_l)b_{h,t} + \omega_m b_{m,t} + \omega_l b_{l,t} \tag{14}$$

$$y_t = \omega_l c_{l,t} + \omega_m c_{m,t} + (1 - \omega_m - \omega_l)c_{h,t} + AC_t + NX_t \tag{15}$$

$$AC_t = \frac{\psi}{2}(b_{h,t} - \bar{b}_h)^2 + \frac{\psi(z_m)}{2}(b_{m,t} - \bar{b}_m)^2 \tag{16}$$

$$NX_t = R_t B_{t-1} - B_t \tag{17}$$

[8] If we assume that income share is a continuous function of population share the Gini coefficient can be written as a function of income and population shares: $Gini = 1 - 2[.5z_l\omega_l + (z_l + .5z_m)\omega_m + (1 - .5z_h)\omega_h]$.

3.6 Consumption, credit shocks and inequality

We are interested in the model's prediction for how aggregate consumption changes in the face of credit supply shocks, here given by $\epsilon_{\mu,t}$. We do not calibrate the model, *per se*, but assume plausible values for all parameters and then compute the dynamic response of consumption to an identical credit shock for alternative values of the income share of the middle-earner group. In the baseline model, which we refer to as the low-inequality regime, middle-earner income share is set at $z_m = 0.55$, which we assume is sufficiently high to eliminate the credit constraint for this group. Alternatively, in the high-inequality regime we set $z_m = 0.45$ and assume that the credit constraint binds. We then solve the model given all parameter values and compute the impulse responses of consumption to credit shocks for each case.

Table 4 reports the parameter values for the alternative parameterizations, along with the implied Gini coefficients. Note that as the middle-earner income share falls from 0.55 to 0.45, the Gini coefficient rises from 0.20 to 0.28. An increase in the Gini coefficient of 0.08 is slightly larger than the standard deviation of 0.068 we observe in the data. Given the population shares, in the low-inequality regime 20% of the population is credit constrained; in the high-inequality regime that magnitude rises to 80%.

Figure 6 shows the model's impulse response functions for consumption under the two cases, along with the (exogenous) response of the loan-to-income ratio μ_t (red curve), which is the same for both scenarios. The shock in this case is a one-time positive impulse to $\epsilon_{\mu,t}$ equal to its assumed standard deviation of 0.02. The shock yields a persistent, but not permanent, effect on μ_t. The general dynamic patterns for consumption are similar to what is seen in the data – a persistent response and a short-run increase in consumption with a declining effect over time.[9] Also evident is the variation in magnitudes given the change in income shares. For the low-inequality country, consumption exhibits an immediate 0.04%

[9] We interpret period 1 in the model to the second quarter horizon in the data, given the Cholesky ordering in the VAR.

rise which essentially falls to zero after four quarters. On the other hand the high-inequality country experiences a 0.22% increase in consumption on impact – more than five times the low-inequality effect – as well as a larger decline after one year.

Although we do no perform a precise calibration exercise, the model simulations accord well with the data. And while we cannot make direct comparisons with Figure 1, note that the countries with large Gini coefficients – Argentina, Brazil, Mexico, Russia and Turkey (each of these countries has a Gini coefficient over .40, well above the average of .33) – have estimated impulse response functions similar to the model's high-inequality country, while the Czech Republic, Netherlands, and other countries with lower Gini coefficients behave more like the low-inequality country.

4 Conclusion

The aim of this paper is to examine the link between income inequality and the 'household demand channel' of credit supply shocks. We document such a link in the data for a sample of developed and developing countries. Time series models for each country suggest that credit shocks temporarily raise aggregate consumption (which is consistent with the extant literature), while cross-country regressions using the time-series estimates show that the short-run rise and long-run decline of spending is larger for countries with a more unequal distribution of income than other countries. We also provide a theoretical model that associates inequality with the extent of binding credit constraints to illustrate a potential mechanism that can drive the empirical results. Simulations show the model generates consumption responses similar to those estimated from the data.

The empirical link that we have documented here between inequality and the nature of credit shocks is important. However, we see this paper as taking only the first steps toward a more complete and informative analysis. Most importantly, we have not provided a tight link between the data and the theoretical model, which is needed to better understand the extent to which the model's mechanisms are relevant for generating the data. We leave this essential analysis for future work.

Appendix

We estimate structural impulse response functions in two ways: standard inversion of a VAR model and the local projections approach of Jorda (2005). In general, for n-dimensional vector process Y_t, the set of impulse response functions for forecast horizon k is given by

$$IRF(k) = E(Y_{t+k}|Y_t + dY_t, Y_{t-1}, Y_{t-2}, \cdots) - E(Y_{t+k}|Y_t, Y_{t-1}, Y_{t-2}, \cdots), \tag{18}$$

where dY_t is taken to be an unexpected innovation or shock. $IRF(k)$ is an $n \times 1$ vector that measures the change in the conditional projection of Y_{t+k} given an impulse in the vector Y_t, and is independent of the data generating process. In our application,

$$Y_t = \begin{bmatrix} c_t \\ D_t^{fy} \\ D_t^{hy} \end{bmatrix}. \tag{19}$$

Under the assumption that the data are generated by a linear vector autoregression (assuming a first order model without loss of generality)

$$Y_t = \phi Y_{t-1} + \epsilon_t, \tag{20}$$

where ϵ_t is a vector of reduced-form forecast errors with covariance matrix $E\epsilon_t\epsilon_t' = \Sigma$. Since $dY_t = d\epsilon_t$,

$$IRF(k) = \phi^k d\epsilon_t. \tag{21}$$

We assume the reduced form errors are linear combinations of orthogonal structural shocks, $\epsilon_t = D_0\Omega^{\frac{1}{2}}u_t$ where D_0 is an $n \times n$ matrix of structural parameters with ones along the diagonal and $Eu_tu_t' = \Omega$, a diagonal matrix. This implies that the conditional expectation of Y_{t+k} is updated in response to structural shocks according to

$$IRF(k) = \phi^k D_0 \Omega^{\frac{1}{2}} du_t. \tag{22}$$

This expression can be calculated by estimating ϕ from the VAR in (20) using standard techniques, assuming that D_0 is lower triangular, and identifying $D_0\Omega^{\frac{1}{2}}$ as the Cholesky factor of Σ. The recursive restrictions on D_0 are consistent with our structural interpretation in the text. As we note in the text, we report dynamic responses to unit-valued structural shocks; i.e., the responses are based on D_0 rather than $D_0\Omega^{\frac{1}{2}}$.

Under the local projections approach, we make no assumptions about the data generating process but rely on a more general formulation of linear projection:

$$E(Y_{t+k}|Y_t, \cdots, Y_{t-p}) = B_k Y_t + \gamma_1 Y_{t-1} + \cdots + \gamma_{p_k} Y_{t-p_k} \tag{23}$$

which implies

$$IRF(k) = B_k dY_t = B_k D_0 \Omega^{\frac{1}{2}} du_t \tag{24}$$

where we have made the same mapping from structure to reduced form as above. The only difference between the VAR and local projections approach to estimating structural impulse response functions is how the weighting matrices ϕ_k and B_k are estimated. As we've assumed that D_0 is lower triangular, the local projections estimates of the response of consumption to household credit shocks are the coefficients on the household credit to GDP ratio in the consumption local projections equations.

References

Abdallah, Chadi S. and William D. Lastrapes. 2012. "Home equity lending and retail spending: Evidence from a natural experiment in Texas." *American Economic Journal: Macroeconomics* 4 (4):94–125. URL http://www.aeaweb.org/articles?id=10.1257/mac.4.4.94.

Aizenman, Joshua, Yothin Jinjarak, Hien Thi Kim Nguyen, and Donghyun Park. 2019. "Fiscal space and government-spending and tax-rate cyclicality patterns: A cross-country comparison, 1960–2016." *Journal of Macroeconomics* 60:229–252.

Alter, Adrian, Alan Xiaochen Feng, and Nico Valckx. 2018. "Understanding the macro-financial effects of household debt: A global perspective." *IMF working paper* (WP/18/76).

Bahadir, Berrak and Inci Gumus. 2016. "Credit decomposition and business cycles in emerging market economies." *Journal of International Economics* 103:250 – 262. URL http://www.sciencedirect.com/science/article/pii/S0022199616301167.

Barnichon, Regis and Christian Brownlees. 2019. "Impulse response estimation by smooth local projections." *The Review of Economics and Statistics* 101 (3):522–530. URL https://doi.org/10.1162/rest_a_00778.

Becker, Gary S. and Casey B. Mulligan. 1997. "The endogenous determination of time preference." *The Quarterly Journal of Economics* 112 (3):729–758. URL https://academic.oup.com/qje/article-abstract/112/3/729/1926899.

Cecchetti, Stephen G. 1999. "Legal structure, financial structure, and the monetary policy transmission mechanism." *Economic Policy Review* 5 (2):9 – 28.

Cloyne, James, Kilian Huber, Ethan Ilzetzki, and Henrik Kleven. 2019. "The effect of house prices on household borrowing: A new approach." *American Economic Review* 109 (6):2104–36. URL http://www.aeaweb.org/articles?id=10.1257/aer.20180086.

Farhi, Emmanuel and Iván Werning. 2016. "A theory of macroprudential policies in the presence of nominal rigidities." *Econometrica* 84 (5):1645–1704. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA11883.

Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh. 2017. "The Macroeconomic effects of housing Wealth, housing finance, and limited risk sharing in general equilibrium." *Journal of Political Economy* 125 (1):140–223. URL https://doi.org/10.1086/689606.

Frederick, Shane, George Loewenstein, and Ted O'donoghue. 2002. "Time discounting and time preference: A critical review." *Journal of Economic Literature* 40 (2):351–401. URL https://www.aeaweb.org/articles?id=10.1257/002205102320161311.

Gennaioli, Nicola, Andrei Shleifer, and Robert W. Vishny. 2012. "Neglected risks, financial innovation and financial fragility." *Journal of Financial Economics* 104 (3):452–468.

Guerrieri, Luca and Matteo Iacoviello. 2017. "Collateral constraints and macroeconomic asymmetries." *Journal of Monetary Economics* 90:28–49. URL https://www.sciencedirect.com/science/article/abs/pii/S0304393217300648.

Iacoviello, Matteo. 2005. "House prices, borrowing constraints, and monetary policy in the business cycle." *American Economic Review* 95 (3):739–764. URL https://www.aeaweb.org/articles?id=10.1257/0002828054201477.

———. 2008. "Household debt and income inequality, 1963-2003." *Journal of Money, Credit and Banking* 40 (5):929–965. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1538-4616.2008.00142.x.

Jorda, Oscar. 2005. "Estimation and inference of impulse responses by local projections." *American Economic Review* 95 (1):161–182. URL https://www.aeaweb.org/articles?id=10.1257/0002828053828518.

Justinian, A., G.E. Primiceri, and A. Tambalotti. 2015. "Credit supply and the housing boom." *National Bureau of Economic Research* w20874.

Korinek, Anton and Alp Simsek. 2016. "Liquidity trap and excessive leverage." *American Economic Review* 106 (3):699–738. URL http://www.aeaweb.org/articles?id=10.1257/aer.20140289.

Kumhof, Michael, Romain Ranciere, and Pablo Winant. 2015. "Inequality, leverage, and crises." *American Economic Review* 105 (3):1217–45. URL http://www.aeaweb.org/articles?id=10.1257/aer.20110683.

Lastrapes, William D. and W. Douglas McMillin. 2004. "Cross-Country variation in the liquidity effect: The role of financial markets." *The Economic Journal* 114 (498):890–915. URL https://dx.doi.org/10.1111/j.1468-0297.2004.00248.x.

Ludvigson, Sydney. 1999. "Consumption and credit: a model of time-varying liquidity constraints." *Review of Economics and Statistics* 81 (3):434–447. URL https://www.mitpressjournals.org/doi/abs/10.1162/003465399558364.

Mian, Atif and Amir Sufi. 2018. "Finance and business cycles: The credit-driven household demand channel." *Journal of Economic Perspectives* 32 (3):31–58. URL http://www.aeaweb.org/articles?id=10.1257/jep.32.3.31.

Mian, Atif, Amir Sufi, and Emil Verner. 2017. "Household debt and business cycles worldwide." *The Quarterly Journal of Economics* 132 (4):1755–1817. URL https://doi.org/10.1093/qje/qjx017.

Mian, Atif R., Amir Sufi, and Emil Verner. 2019. "How does credit supply expansion affect the real economy? The productive capacity and household demand channels." *Working Paper* .

Mumtaz, Haroon, Gabor Pinter, and Konstantinos Theodoridis. 2018. "What do VARs tell US about the impact of a credit supply shock?" *International Economic Review* 59 (2):625–646. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/iere.12282.

Schmitt-Grohé, Stephanie and Martin Uribe. 2016. "Downward nominal wage rigidity, currency pegs, and involuntary unemployment." *Journal of Political Economy* 124 (5):1466 – 1514. URL https://EconPapers.repec.org/RePEc:ucp:jpolec:doi:10.1086/688175.

Solt, Frederick. 2016. "The standardized world income inequality database*." *Social Science Quarterly* 97 (5):1267–1281. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/ssqu.12295.

Svirydzenka, Katsiaryna. 2016. "Introducing a new, broad-based index of financial development." *IMF working paper* (WP/16/5).

Table 1: Summary Statistics

No.	*Country*	*Begin*	*End*	c	$\frac{hhdebt}{gdp}$	$\frac{firmdebt}{gdp}$	fd_{index}	$gini_{index}$
1	Argentina	1994:1	2017:1	11.01	0.048	0.242	0.337	0.433
2	Australia	1991:1	2017:2	25.82	0.848	0.691	0.768	0.320
3	Austria	1996:1	2017:3	6.54	0.487	0.845	0.628	0.272
4	Belgium	1995:1	2017:3	6.57	0.464	1.187	0.534	0.256
5	Brazil	1997:1	2017:3	10.49	0.161	0.409	0.462	0.497
6	Canada	1990:1	2017:3	7.58	0.716	0.901	0.748	0.305
7	Czech Republic	1995:4	2017:3	8.38	0.191	0.586	0.354	0.246
8	Denmark	1995:1	2017:3	7.82	1.074	0.894	0.647	0.238
9	Finland	1990:1	2017:3	5.32	0.461	0.977	0.561	0.243
10	France	1990:1	2017:1	7.81	0.418	1.032	0.671	0.287
11	Germany	1991:1	2017:3	8.15	0.611	0.547	0.717	0.273
12	Greece	1995:1	2017:3	6.10	0.374	0.497	0.513	0.336
13	Hungary	1995:1	2016:4	9.97	0.197	0.659	0.421	0.280
14	Israel	1992:3	2017:3	6.93	0.375	0.743	0.517	0.348
15	Italy	1995:1	2016:4	7.75	0.319	0.672	0.680	0.328
16	Japan	1994:1	2016:4	13.47	0.643	1.148	0.733	0.302
17	Korea	1991:1	2017:3	13.79	0.631	0.933	0.735	0.292
18	Mexico	1995:1	2017:3	20.94	0.117	0.199	0.349	0.466
19	Netherlands	1995:1	2017:3	7.59	0.993	1.204	0.725	0.261
20	New Zealand	1998:2	2017:3	8.12	0.798	0.859	0.545	0.326
21	Norway	1990:1	2017:1	7.67	0.692	1.177	0.620	0.246
22	Poland	1996:1	2017:3	7.73	0.204	0.353	0.394	0.306
23	Portugal	1995:1	2017:3	5.81	0.696	1.035	0.613	0.341
24	Russia	1998:1	2017:3	9.45	0.080	0.354	0.358	0.402
25	Singapore	1991:1	2017:2	5.45	0.413	0.863	0.680	0.390
26	Spain	1995:1	2017:3	6.88	0.613	0.947	0.741	0.327
27	Sweden	1990:1	2017:2	8.12	0.607	1.172	0.658	0.243
28	Switzerland	1999:4	2017:1	7.12	1.118	0.931	0.896	0.295
29	Thailand	1993:1	2017:3	8.92	0.495	1.175	0.535	0.429
30	Turkey	1990:1	2017:3	7.29	0.073	0.301	0.394	0.425
31	UK	1990:1	2016:2	7.79	0.742	0.780	0.798	0.338
32	US	1990:1	2017:3	9.04	0.768	0.634	0.806	0.365
	μ			9.11	0.513	0.780	0.598	0.325
	σ			4.20	0.289	0.297	0.153	0.068

Notes: c, $\frac{hhdebt}{gdp}$, $\frac{firmdebt}{gdp}$, fd_{index} and $gini_{index}$ indicate respectively averages of log real consumption (in local currency), household debt to GDP, firm debt to GDP, financial development and gini index for each country over the sample period. μ and σ are the cross-sectional mean and standard deviation (across the sample of 32 countries) of the reported variables.

Table 2: Summary consumption response to unit and std. deviation shocks

No.	*Country*	c_{max}		c_{q2}		c_{q4}		c_{q8}	
		unit	std. dev	unit	std. dev	unit	std. dev	unit	std. dev
1	Argentina	0.068	0.010	0.062	0.009	0.053	0.008	-0.004	-0.001
2	Australia	0.005	0.004	0.001	0.001	0.003	0.002	0.004	0.003
3	Austria	0.006	0.003	0.002	0.001	0.006	0.003	0.005	0.002
4	Belgium	0.000	0.000	-0.001	0.000	0.000	0.000	-0.001	0.000
5	Brazil	0.016	0.007	0.012	0.005	0.014	0.006	0.014	0.006
6	Canada	0.000	0.000	-0.001	0.000	0.000	0.000	-0.002	-0.001
7	Czech Republic	0.003	0.002	0.003	0.001	0.002	0.001	-0.001	-0.001
8	Denmark	0.003	0.003	0.001	0.001	0.002	0.002	0.002	0.002
9	Finland	0.003	0.001	0.002	0.001	0.003	0.001	-0.002	-0.001
10	France	0.004	0.001	-0.002	-0.001	0.004	0.001	0.002	0.001
11	Germany	0.003	0.001	-0.002	-0.001	0.001	0.000	0.002	0.001
12	Greece	0.017	0.005	0.007	0.002	0.011	0.003	0.016	0.005
13	Hungary	0.003	0.002	0.003	0.002	-0.001	-0.001	-0.005	-0.003
14	Israel	0.002	0.001	-0.003	-0.001	-0.006	-0.003	-0.002	-0.001
15	Italy	0.024	0.004	0.001	0.000	0.016	0.003	0.022	0.004
16	Japan	0.001	0.001	0.000	0.000	-0.002	-0.001	0.000	0.000
17	Korea	0.000	0.000	-0.004	-0.003	-0.010	-0.008	-0.018	-0.013
18	Mexico	0.026	0.004	0.020	0.003	0.025	0.004	0.001	0.000
19	Netherlands	0.002	0.002	0.000	0.000	0.002	0.001	0.002	0.001
20	New Zealand	0.011	0.004	0.005	0.002	0.010	0.004	0.009	0.003
21	Norway	0.004	0.003	0.003	0.002	0.004	0.003	0.001	0.001
22	Poland	0.000	0.000	-0.002	-0.001	-0.001	-0.001	-0.006	-0.002
23	Portugal	0.019	0.008	0.004	0.002	0.012	0.005	0.018	0.008
24	Russia	0.023	0.005	0.022	0.005	0.023	0.005	-0.025	-0.005
25	Singapore	0.004	0.002	-0.004	-0.002	0.002	0.001	0.003	0.002
26	Spain	0.011	0.006	0.002	0.001	0.008	0.004	0.010	0.006
27	Sweden	0.009	0.003	0.002	0.001	0.008	0.002	0.008	0.002
28	Switzerland	0.002	0.001	-0.003	-0.001	-0.003	-0.001	-0.001	0.000
29	Thailand	0.010	0.005	0.006	0.003	0.010	0.005	0.003	0.002
30	Turkey	0.022	0.006	0.022	0.006	0.006	0.001	-0.001	0.000
31	UK	0.008	0.003	0.000	0.000	0.007	0.002	0.008	0.002
32	US	0.006	0.003	0.002	0.001	0.004	0.002	0.006	0.003
	μ	0.010	0.003	0.005	0.001	0.007	0.002	0.002	0.001

Notes: This table reports sample summary consumption responses from recursive VAR to unit and std. deviation shocks. c_{max} indicates peak consumption response over 12 quarters, and c_{qj} indicates consumption response at quarter j. The last row reports the cross-sectional mean response across the sample of 32 countries at the chosen horizon.

Table 3: Cross section, VAR, unit and std. deviation shocks

y	β_0		β_1		β_2		R^2	
	unit	std. dev	unit	std. dev	unit	std. dev	unit	std. dev
c_{max}	-0.005	-0.002	-0.025	-0.002	0.092	0.021	0.43	0.42
	(0.577)	(0.236)	(0.064)	(0.274)	(0.002)	(0.000)		
c_{q2}	0.003	0.001	-0.035	-0.008	0.072	0.014	0.48	0.56
	(0.763)	(0.542)	(0.006)	(0.000)	(0.008)	(0.001)		
c_{q4}	-0.004	-0.002	-0.022	-0.003	0.073	0.017	0.40	0.26
	(0.606)	(0.409)	(0.051)	(0.188)	(0.002)	(0.000)		
c_{q8}	-0.018	-0.007	0.020	0.006	0.025	0.013	0.10	0.07
	(0.009)	(0.008)	(0.037)	(0.098)	(0.110)	(0.026)		
c_{q12}	-0.025	-0.010	0.043	0.012	-0.010	0.006	0.17	0.13
	(0.048)	(0.007)	(0.042)	(0.025)	(0.779)	(0.469)		
c_{q16}	-0.025	-0.010	0.047	0.013	-0.024	0.002	0.19	0.17
	(0.066)	(0.009)	(0.041)	(0.014)	(0.518)	(0.836)		
c_{q20}	-0.021	-0.009	0.039	0.011	-0.023	0.001	0.21	0.19
	(0.046)	(0.007)	(0.032)	(0.008)	(0.463)	(0.916)		
c_{q24}	-0.016	-0.007	0.025	0.008	-0.014	0.001	0.20	0.16
	(0.033)	(0.011)	(0.036)	(0.015)	(0.490)	(0.860)		
$\sum_{q1}^{q12} c$	-0.150	-0.059	0.080	0.035	0.405	0.146	0.11	0.08
	(0.009)	(0.018)	(0.367)	(0.286)	(0.002)	(0.006)		
$\sum_{q12}^{q24} c$	-0.290	-0.115	0.522	0.147	-0.259	0.029	0.20	0.18
	(0.049)	(0.006)	(0.036)	(0.012)	(0.532)	(0.798)		
$\sum_{q1}^{q24} c$	-0.415	-0.165	0.559	0.169	0.156	0.168	0.13	0.11
	(0.019)	(0.006)	(0.062)	(0.035)	(0.744)	(0.248)		
$(\bar{c})_{q1}^{q12} - (\bar{c})_{q12}^{q24}$	0.010	0.004	-0.033	-0.008	0.054	0.010	0.32	0.31
	(0.276)	(0.064)	(0.022)	(0.005)	(0.040)	(0.134)		
$(\bar{c})_{q1}^{q8} - (\bar{c})_{q16}^{q24}$	0.014	0.006	-0.046	-0.011	0.071	0.012	0.34	0.32
	(0.242)	(0.043)	(0.017)	(0.003)	(0.042)	(0.146)		

Notes: The regression is $y_i = \beta_0 + \beta_1 z_{i1} + \beta_2 z_{i2} + \epsilon_i$, where y_i is a summary measure of each country's dynamic response of consumption to household credit shocks in country i; z_{i1} and z_{i2} are nation-wide financial development and gini index. c_{max} indicates the peak response over 12 quarters. c_{qj}, $\sum_{qj}^{qk} c$, and $\left[(\bar{c})_{qj}^{qk} - (\bar{c})_{ql}^{qm}\right]$ indicate respectively response at quarter j, cumulative response over quarters j to k, and average response over quarters j to k less l to m. The first column reports coefficient responses to a unit shock, the second to a standard deviation shock.

Table 4: Parameter values for low- and high-inequality regimes. Values below the dashed line are determined by values above the dashed line.

Parameter	Inequality Low	Inequality High	Description
β_l	0.96	0.96	Discount factor: low earners
β_m	0.99	0.96	Discount factor: middle earners
β_h	0.99	0.99	Discount factor: high earners
σ	1.00	1.00	Relative risk aversion
ω_l	0.20	0.20	Population share: low earners
ω_m	0.60	0.60	Population share: middle earners
$\bar{R}$	1.01	1.01	Real interest rate
θ	0.60	0.60	Inertia in borrowing limit
ψ	0.006	0.006	Adjustment cost parameter
ρ_m	0.90	0.90	Persistence of credit shock
σ_m	0.02	0.02	Standard deviation of credit shock
$\bar{\mu}_l$	0.50	0.50	Loan-to-income: low earners
$\bar{\mu}_m$	–	0.50	Loan-to-income: middle earners
$\bar{\mu}_h$	0.50	0.50	Loan-to-income: high earners
$\bar{b}_m$	0.50	–	Steady-state borrowing: middle earners
$\bar{b}_h$	0.50	0.50	Steady-state borrowing: high earners
z_l	0.10	0.10	Income share: low earners
z_m	**0.55**	**0.45**	Income share: middle earners
z_h	**0.35**	**0.45**	Income share: high earners
b_l/b	0.11	0.13	Borrowing share: low earners
b_m/b	0.67	0.61	Borrowing share: middle earners
B_t/y_t	0.46	0.38	Aggregate borrowing to output
Gini	0.20	0.28	Gini coefficient: income distribution

Figure 1: Impulse responses for real consumption to a unit household credit shock from a recursive VAR

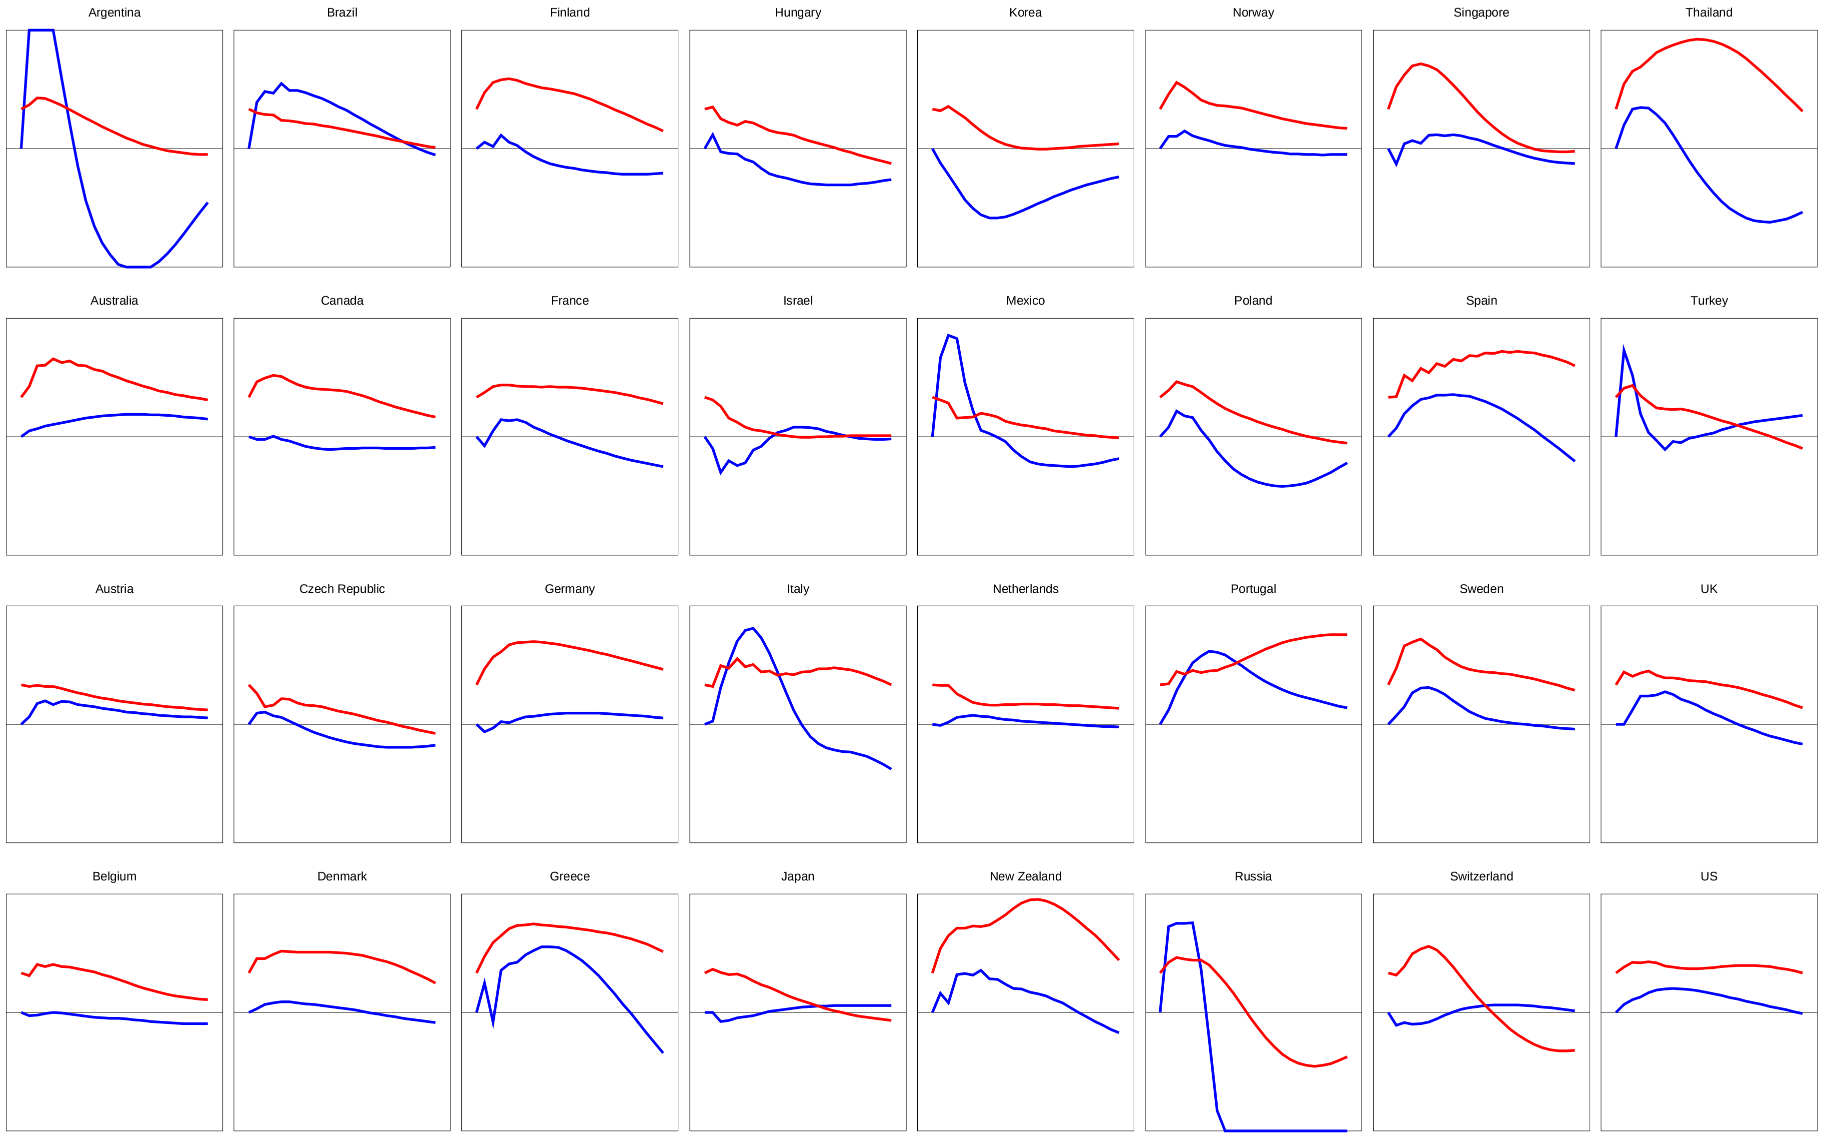

Notes: This figure presents the impulse responses for real consumption to a unit household credit shock from a recursive VAR in log real consumption, firm debt to GDP, and household debt to GDP. The solid blue line plots the consumption response while the solid red line the shock. Both lines are plotted against the left vertical scale with maximum and minimum of .03 and -.03.

Figure 2: Impulse responses for real consumption to a std. deviation household credit shock from a recursive VAR

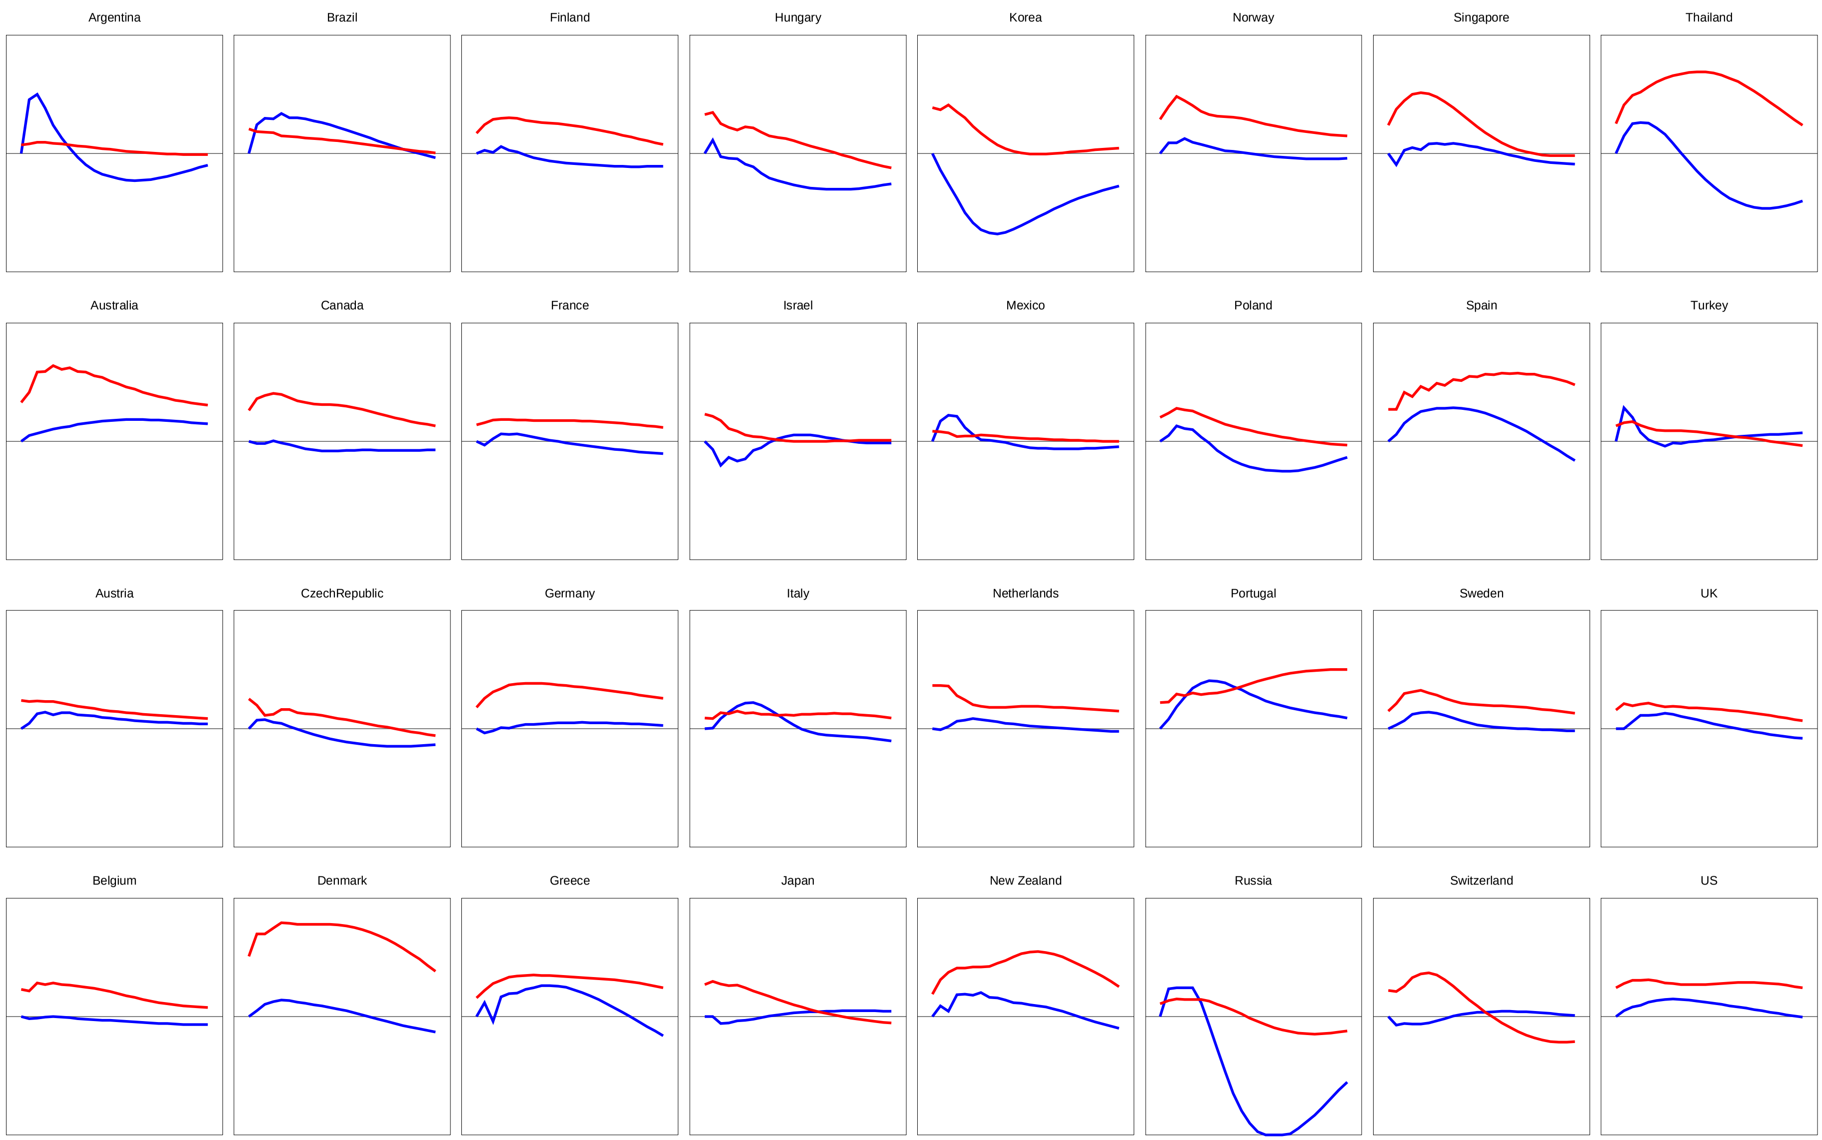

Notes: This figure presents the impulse responses for real consumption to a one standard deviation household credit shock from a recursive VAR in log real consumption, firm debt to GDP, and household debt to GDP. The solid blue line plots the consumption response while the solid red line the shock. Both lines are plotted against the left vertical scale with maximum and minimum of .02 and -.02.

Figure 3: Summary consumption responses from recursive VAR

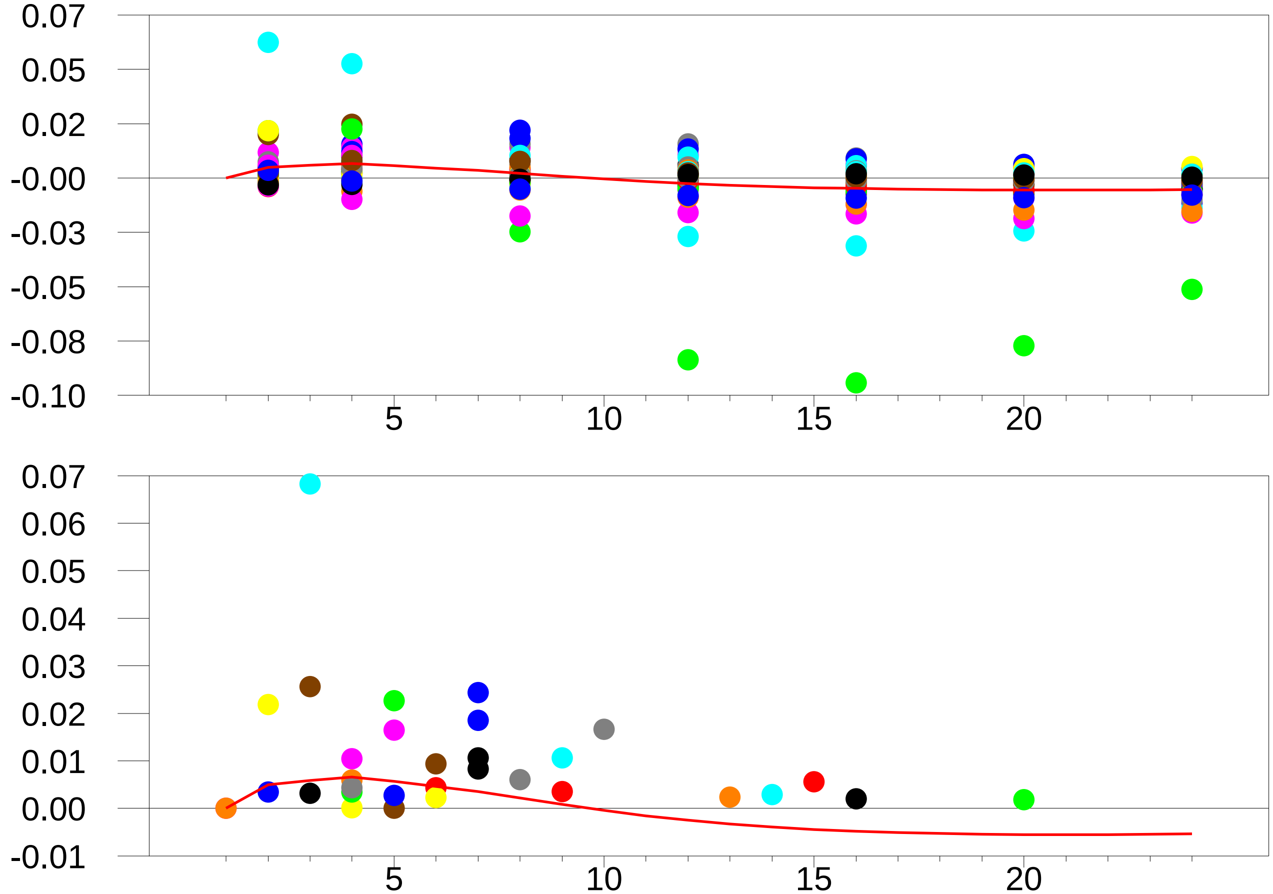

Notes: This figure summarizes the dynamic responses of consumption to the unit household credit shock from the recursive VAR. The top panel plots the responses at horizons 2, 4, 8, 12, 16, 20 and 24 quarters. The solid red curve plots the cross-sectional mean response across the sample of 32 countries for each horizon from 1 to 24. The bottom panel plots each country's peak response, at the horizon at which that peak occurs (with the cross-sectional mean superimposed).

Figure 4: Impulse responses for real consumption to a unit household credit shock from Jorda local projections

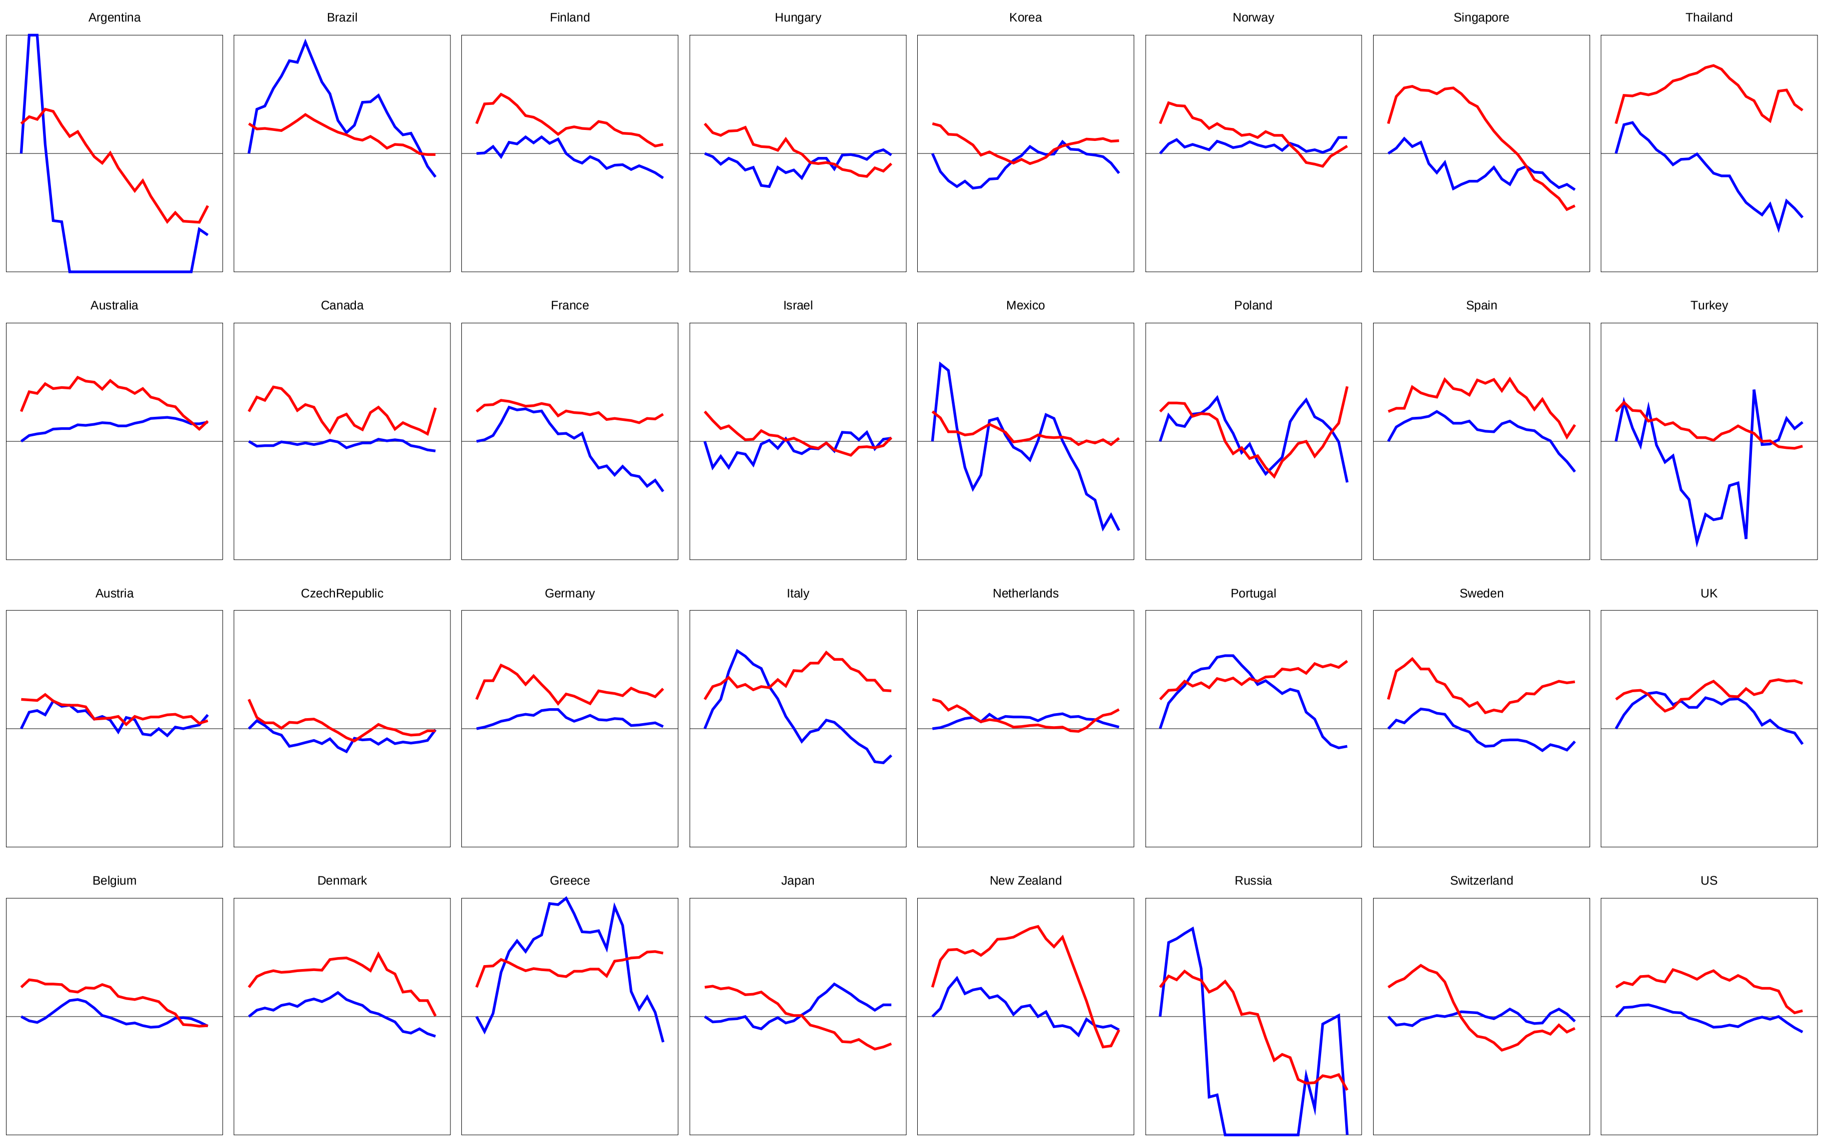

Notes: This figure presents the impulse responses for real consumption to a unit household credit shock from Jorda local projections in log real consumption, firm debt to GDP, and household debt to GDP. The solid blue line plots the consumption response while the solid red line the shock. Both lines are plotted against the left vertical scale with maximum and minimum of .04 and -.04.

Figure 5: Summary consumption responses from Jorda local projections

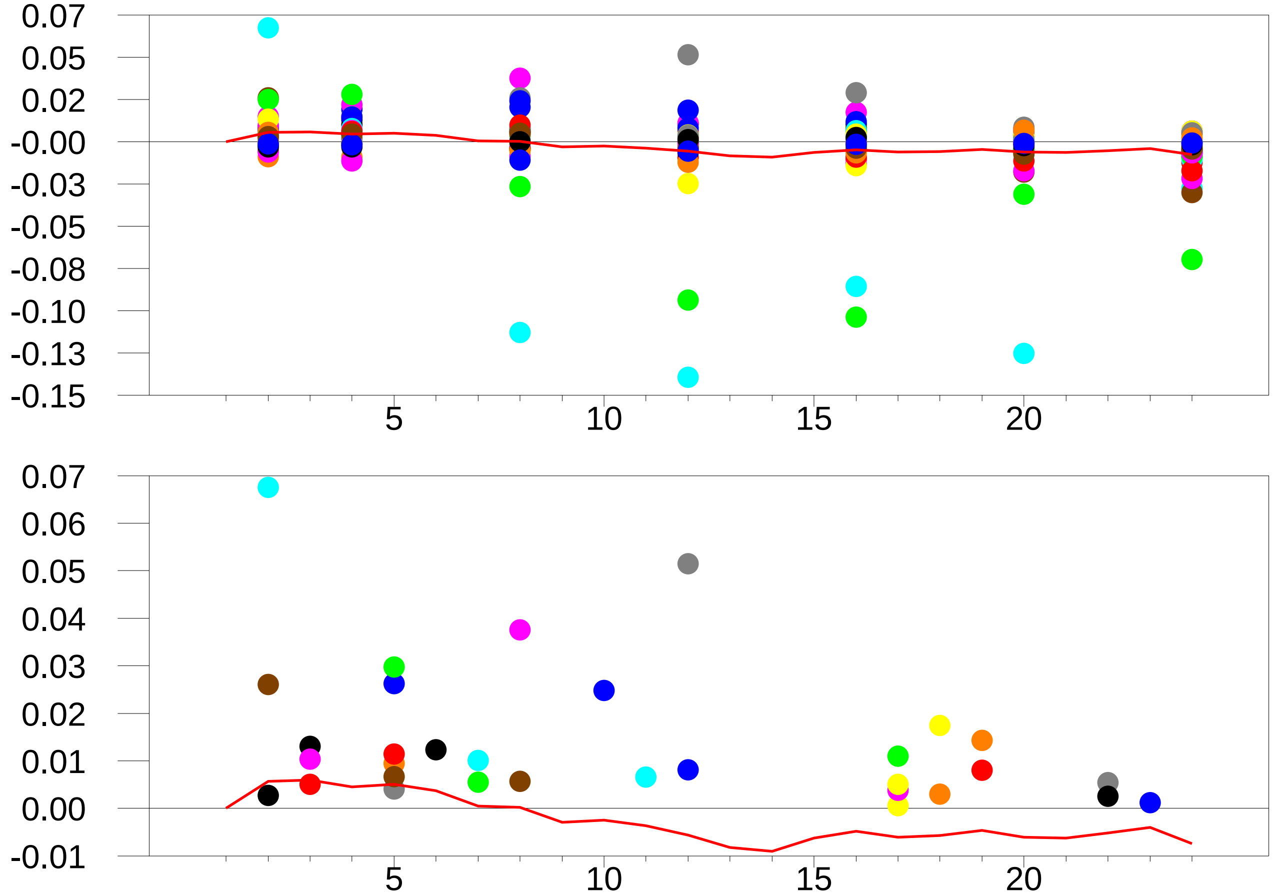

Notes: This figure summarizes the dynamic responses of consumption to the unit household credit shock from the local projections. The top panel plots the responses at horizons 2, 4, 8, 12, 16, 20 and 24 quarters. The solid red curve plots the cross-sectional mean response across the sample of 32 countries for each horizon from 1 to 24. The bottom panel plots each country's peak response, at the horizon at which that peak occurs (with the cross-sectional mean superimposed).

Figure 6: Impulse responses of consumption to a household credit shock from theoretical framework

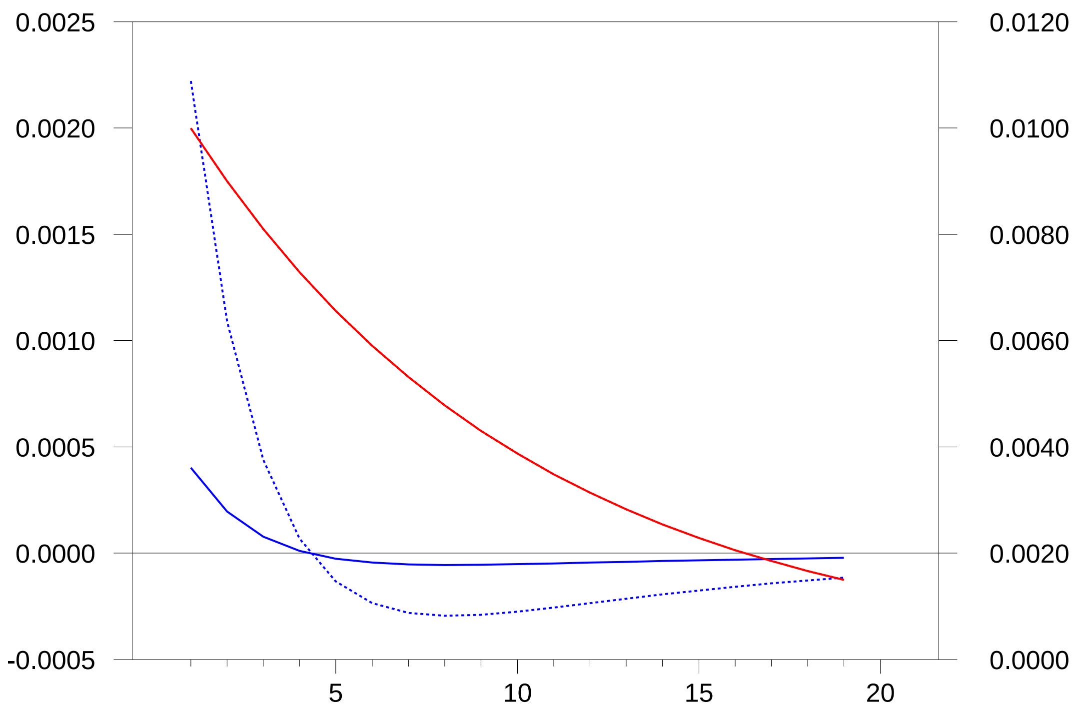

Notes: This figure reports the results from our theoretical framework. The solid red line (corresponding to the right scale) plots the shock to loan-to-income ratio, μ; while the dashed and solid blue lines (corresponding to the left scale) plot the responses of consumption to the shock in high and low income countries respectively.