

Time-Varying Influence of Household Debt on Inequality in United Kingdom

Edmond Berisha*
Feliciano School of Business
Montclair State University
Montclair, NJ 07043, USA
Email: berishae@mail.montclair.edu

David Gabauer
Institute of Applied Statistics,
Johannes Kepler University,
Altenbergerstraße 69, 4040 Linz, Austria;
Department of Business and Management,
Webster Vienna Private University,
Praterstraße 23, 1020 Vienna, Austria.
Email: david.gabauer@hotmail.com

Rangan Gupta
Department of Economics
University of Pretoria
Pretoria, South Africa
Email: rangan.gupta@up.ac.za

Chi Keung Marco Lau
Huddersfield Business School,
University of Huddersfield,
Huddersfield, HD1 3DH, UK
Email: c.lau@hud.ac.uk

* Corresponding author.

Abstract

The United Kingdom (UK) in terms of income inequality is ranked among the highest in Europe. Likewise, within the last four decades, UK is characterized with drastic increases in household debt. In this paper, we analyze time-varying predictability of growth in household debt for growth in income (and consumption) inequality based on a high-frequency (quarterly) data set over 1975:Q2 to 2016:Q1. Results indicate that the growth in household debt has a strong predictive power, both for within and out-of-samples, on growth rate of income (and consumption) inequality in the UK. Interestingly, the strength of the predictive power is found to have increased after 2008. Based on time-varying impulse response functions, we also find that higher growth rate in household debt corresponds with subsequent increases in income inequality.

JEL Code: C32, C53, D63, E30, E40, G51, R31

Keywords: Household Debt; Inequality; Time-Varying Predictions

I. Introduction

Recent theoretical and empirical macroeconomic research have identified household debt as a significant contributor in shaping business cycles. Jorda et al (2013) claim that the periods after credit booms are associated with slower growth, investment spending and credit growth than usual. Similarly, Mian et al (2017) identify the credit supply shocks as an essential driver of economic fluctuations. They find that increases in the household debt to GDP ratio over a three year period in a given country predicts subsequently lower output growth and higher unemployment. An interesting remark is that the observed credit booms from the data was mainly concentrated among the households outside the top of the income distribution (Cynamon and Fazzari (2013), Barba and Pivetti (2008)). Thus, many economists have emphasized the role of household debt in driving redistribution of income in a country. The mechanism is that increases in the size and importance of the financial sector across the developed and developing economies over the past few decades have made it easier for high-income households to back their financial investments with loans to the rest of the population. As such, lower income households were able to sustain their consumption levels by borrowing; whereas, the high income households gained additional income by providing savings later used as loans, further exacerbating income inequality. Correspondingly, Saez (2016) suggests that large increases in debt for households in the bottom 90% of the income distribution implies that these households have been saving 0% of their income over the last 30 years. However, the top wealth holders have saved significantly, in part because their incomes have increased so much that they can afford to save large shares of their incomes. The result is a large increase in income and wealth inequality that is likely to persist.

We contribute to the existing literature by providing a sharper picture of the role of household debt in driving redistribution of income. Specifically, we explore the predictive power of the growth of household debt for growth in income (and consumption) inequality in the UK. Over the last four decades, UK has experienced dramatic increases in income inequality (Mumtaz and Theophilopoulou (2017)) and is considered an outlier of extreme inequality in the European context (Dorling, 2015). In the Appendix, Figure A1(a) shows the time series of household debt and income inequality (Gini coefficient), and Figure A1(b) shows the time series of household debt and consumption inequality (Gini coefficient) in the UK. From both figures we observe a growth in massive inequality, with income (consumption) inequality growth between March, 1975 to March, 2016 being 13.63% (11.19%) in our data set. Likewise, we also observe drastic increases in household debt from 31% of to 85.5 % of Gross Domestic Product (GDP). Another underlying reason why we focus on the UK is that income (and consumption) inequality data are available at a high frequency, i.e., on a quarterly basis for over 40 years (March, 1975 to March, 2016). This is important since accurate prediction of inequality at a higher frequency should be more relevant to policymakers than at the lower annual frequency, as considered traditionally in the literature involving inequality, including the most recent UK-based study by Wood (2019). The author examines distributional consequences of household indebtedness in Britain using annual data between 1966 and 2016. He demonstrates that expansion of household debt shows positive and significant effects across four measures of inequality in Britain.

From a methodological perspective, this paper is the first study to analyze time-varying predictability of growth in household debt for growth in income (and consumption) inequality based on a high-frequency data set. In this regard, we use the recently proposed multivariate test of time-varying causality in a vector autoregressive (VAR) framework by Wang and Rossi (forthcoming), which is robust to the presence of instabilities. Note the VAR model includes the additional controls of economic growth, real returns of house prices, and real interest rates. Though the time-varying predictability is the primary focus of our paper, since causality tests are silent about the sign of the

impact (if any), we use time-varying impulse response functions to analyze the effect of a shock to the growth in household debt on the growth of inequality in a time-varying parameter VAR model. Finally, we also examine the out-of-sample forecastability of the growth in income inequality from the variation in the household debt.

To preview, results indicate that the growth in household debt has strong predictive power on growth in income (and consumption) inequality in the UK, though the strength of the effect is indeed time-varying. Moreover, increases in the growth in debt corresponds with subsequent higher growth in income inequality. In addition, over and above the information contained in other control variables, growth in household debt is found to produce forecasting gains over an out-of-sample period for the growth of income inequality. The rest of the paper proceeds as follows. Section 2 discusses the data and the methodology. Section 3 presents the empirical results, and Section 4 concludes.

II. Data and Methodology

1) Data

We use quarterly data from 1975:Q1 to 2016:Q1, based on availability of measures of inequality. The income inequality data is taken from Mumtaz and Theophilopoulou (2017). To construct income inequality measures they use income equivalized by dividing with the square root of the number of people in a household. The inequality measures are computed using survey data on income and consumption from the family expenditure survey (FES).¹ Mumtaz and Theophilopoulou (2017) provide an extensive documentation of the construction of the data and the survey. Note that, while the surveys are recorded at an annual frequency, Mumtaz and Theophilopoulou (2017) assign households to different quarters within a year based on the date of the survey interviews, which, in turn, allows them to calculate the measures of inequality at a quarterly frequency. Note that, these authors remove any households reporting zero or negative income, when constructing the income-based measures of inequality.² In our main analysis, we consider the Gini coefficient of income inequality. However, as part of additional analyses, we also present the results in the involving the Gini coefficient of consumption inequality (for total consumption per capita of a household), as well as, the standard deviation (of the data in natural logs), and the difference between the 90th and 10th percentile (with the data in natural logs) associated with both income and consumption. These five additional measures of inequality is also sourced from Mumtaz and Theophilopoulou (2017).

Household debt defined as credit to households and non-profit institutions serving households (NPISHs) from all sectors at market value as a percentage of GDP. The data source for the household debt is Bank of International Settlements. We consider that development in overall economic conditions, monetary policy, and housing prices might contaminate the relationship between household debt and income (consumption) inequality. Thus, as additional controls are concerned, we use GDP for the UK to capture current economic conditions. Real house prices (computed by deflating the nominal price index with the Consumer Price Index (CPI)) and real interest rate (nominal three-month Treasury bill rate less the CPI-based inflation rate) are used as measures for development in housing prices and monetary policy. The data source used to derive the three controls is the Main Economic Indicators database of the Organisation for Economic Co-operation and Development (OECD).

¹ The data is downloadable from: <https://discover.ukdataservice.ac.uk/series/?sn=200016> and <https://discover.ukdataservice.ac.uk/series/?sn=2000028>.

² We would like to thank Professor Haroon Mumtaz for kindly sharing the inequality data.

Since our econometric approach, which we describe below requires us to work with stationary data, we convert all data into their corresponding growth rates,³ and depict them as: GI_j , $j=1..6$, corresponding to the six measures of income and consumption inequalities (Gini coefficient, the standard deviation, and the difference between the 90th and 10th percentile) respectively; GD: growth in household debt; GRGDP: growth of real GDP; RHR: real housing returns, and; RIR: real interest rate. Due to the transformations, our effective sample starts from 1975:Q2 to 2016:Q1, giving us a total of 164 observations. All variables have been plotted in Figure A1(c) in the Appendix of the paper.

2) Methodology

The simplicity of the classical linear Granger causality test makes it one of the most commonly used methods for testing in-sample predictability. However, VAR model-based analyses face major technical difficulties in handling macro-economic data that are subject to instabilities, thereby causing the associated estimates of VARs to be also sensitive to instabilities (Boivin and Giannoni, 2006; Clark and McCracken, 2006; Rossi, 2013). Moreover, the traditional Granger-causality test requires stationarity of the variables, which in turn may lead to an erroneous inference in the presence of instabilities. In order to overcome these problems, Wang and Rossi (forthcoming) propose a robust causality test, which is more powerful than the traditional Granger-causality test, following the time-varying methodologies suggested by Rossi (2005). Given that we find statistical evidence of regime-changes in our data, we use Wang and Rossi's approach to analyze the time-varying impact of GD on GI_j , $j=1, \dots, 6$, and hence, provide a more appropriate inference of the effect rather than a constant parameter Granger causality method. Besides, understandably, the time-varying approach helps us to depict the time-variation in the strength of predictability.

In this study, we consider the following VAR model with time-varying parameters:

$$y_t = \Psi_{1,t}y_{t-1} + \Psi_{2,t}y_{t-2} + \dots + \Psi_{p,t}y_{t-p} + \varepsilon_t \quad (1)$$

where $\Psi_{j,t}$, $j = 1, \dots, p$ are functions of time varying coefficient matrices, $y_t = [y_{1,t}, y_{2,t}, \dots, y_{n,t}]'$ is an $(n \times 1)$ vector and the idiosyncratic shocks ε_t are assumed to be heteroscedastic and serially correlated.

The variables included in our VAR constitutes of five endogenous variables namely, GI_j , $j=1, \dots, 6$, GD, GRGDP, RHR, and RIR. We test the null hypothesis that GD does not Granger cause GI_j for all t where the null hypothesis is $H_0: \phi_t = 0$ for all $t = 1, 2, \dots, T$, given that ϕ_t is appropriate subset of $vec(\Psi_{1,t}, \Psi_{2,t}, \dots, \Psi_{p,t})$. To this end, Rossi (2005) suggested four alternative test statistics namely: the exponential Wald (*ExpW*), mean Wald (*MeanW*), Nyblom (*Nyblom*) and Quandt Likelihood Ratio (*SupLR*) tests. Based on the Akaike Information Criterion (AIC), the VAR model is estimated with 4 lags. We use an end-point trimming of 15% commonly used in the structural break literature, which in turn amounts us to losing 25 observations from both ends.

III. Empirical Results

In Table 1, to analyze the predictive ability of GD on GI_1 , we first started with the standard constant parameter Granger causality test and found that GD does indeed Granger cause GI_1 at the 5% level of significance. However, based on the powerful *UDmax* and *WDmax* tests of Bai and Perron (2003), used to detect 1 to M structural breaks in the GI_1 equation of the VAR(4) model, allowing for

³ Complete details of the unit root tests are available upon request from the authors.

heterogeneous error distributions across the breaks and 15% trimming, yielded 5 break points at: 1984:Q2, 1990:Q2, 1997:Q1, 2003:Q2, 2009:Q2. Given this evidence of instability, the results from the constant parameter model might not be completely reliable, and hence to obtain reliable inference, we look at the *ExpW*, *MeanW*, *Nyblom*, and *SupLR* tests of Wang and Rossi (forthcoming) based on the time-varying VAR also reported in Table 1. As can be seen, the null of no-Granger causality from GD to GI1 is overwhelmingly rejected at the highest possible level of significance across all the four tests. In other words, the predictive ability of GD for GI1 is in fact time-varying, even though evidence of predictability can also be derived from the constant parameter model.

[INSERT TABLE 1]

Next, in Figure 1, we present the whole sequence of the Wald statistics across time, which gives more information on when the Granger-causality occurs. As can be seen, GD consistently predicts GI1, with the effect following an upward trend since the Global Financial Crisis (GFC). Given the discussion in the literature that growth in inequality has a feedback on the growth of household debt (see for example, Barba and Pivetti (2008), Iacoviello (2008), and Rajan (2010)), in Figure A2 in the Appendix of the paper, we plot the time-varying Wald statistics of GI1 Granger causing GD. A similar picture to the effect of GD on GI1 emerges, with causality consistently running from GI1 to GD and increasing sharply after the GFC.⁴ Furthermore in Figures A3 to A6, we provide the evolution of the Wald statistics over time, whereby we test the time varying causality from GD to GI_{*j*}, *j*=2,..6. As can be seen from the figures, GD consistently causes the alternative measures of income and consumption inequalities over the entire sample period, based on a VAR(4) model with 15% trimming.

[INSERT FIGURE 1]

As indicated above, since 15% trimming leads to a loss of 25 observations from both ends of the sample period, we revisit our results by taking two alternative routes to minimize the loss of observations. In the first case, we re-conduct the Wang and Rossi (forthcoming) test by using a VAR(1) model as suggested by the Schwarz Information Criterion (SIC), which in turn also allows us to use a trimming of only 5% due to the usage of one lag. As can be seen from Figure 2, when we lose 8 observations from both ends, the results are similar to those reported in Figure 1, but the longer coverage of data shows a sudden jump in the time-varying Wald statistics suggesting a sharp increase in predictability towards the end of the sample at around 2013.

[INSERT FIGURE 2]

In the second case, we use the bootstrapped recursive, rolling, and recursive-rolling multivariate tests of Shi et al., (2018, forthcoming) to check for time-varying causality from GD to GI1 using a VAR(1) model. Based on an initial window of 0.20 of the total sample, i.e., 33 quarters and 500 bootstrap replications, and lag chosen by the SIC, as recommended by the original papers, results are reported in Figure 3 for recursive, rolling and recursive-rolling approaches. As can be seen from the figure, we find evidence of causality over the sample period, primarily under the recursive-window estimation, barring the initial few years.

[INSERT FIGURE 3]

⁴ Interestingly, the constant parameter-based Granger causality test could not detect causality running from GI1 to GD even at the 10% level of significance, though *ExpW*, *MeanW*, *Nyblom*, and *SupLR* tests all overwhelmingly rejected the null of no time-varying predictability at the highest level of significance. Complete details of these results are available upon request from the authors.

While the time-varying predictive analysis is the focus of our paper, causality tests are silent about the sign of the impact of GD on GI1. Given this, as a final part of the analysis, we estimate a time-varying parameter VAR model with stochastic volatility (TVP-VAR-SV) as developed by Primiceri (2005), and Del Negro and Primiceri (2015) with 4 lags, and analyze the time-varying impact on GI1 following a shock to GD. Given the evidence of bi-directional causality, we rely on generalized impulse response functions (GIRFs). The TVP-VAR model is estimated using Markov-Chain Monte-Carlo (MCMC) methods with Bayesian inference, based on 50,000 draws after an initial burn-in of 50,000 (i.e., we use a total of 100,000 iterations). The MCMC method assesses the joint posterior distributions of the parameters of interest based on certain prior probability densities that we set in advance, which in turn, are identical to those used in Primiceri (2005), Del Negro and Primiceri (2015). Once the model is estimated, we can produce time-varying impulse response functions of the variables in the model following the one standard deviation of the GD shock. In Figure 4, we present the time-varying response of GI1 over a horizon of 10 quarters. In essence, the impact is sharp and positive on the impact of the GD shock throughout the sample period, with the effect dying down post the second-quarter onwards.

[INSERT FIGURE 4]

Finally, since in-sample predictability does not guarantee out-of-sample predictive gains (Campbell, 2008), we conduct a full-fledged forecasting exercise, by estimating a TVP-VAR models along the lines of Primiceri (2005), Del Negro and Schorfheide (2013), and Del Negro and Primiceri (2015). In particular, we estimate two TVP-VARs of lag order 4, with one containing GI1, GRGDP, RHR, and RIR and the other including GD along with the variables in the first model, which we call them Model 1 and Model 2 respectively. Understandably, if the second model produces lower mean square forecast errors (MSFEs) relative to the first model for GI1, then clearly GD contains out-of-sample predictive information for the future path of the growth in income inequality. Given that our first break date is at 1984:Q2, we design our forecasting experiment to have an in-sample period of 1975:Q2 to 1984:Q1, with the models estimated recursively over the out-sample covering 1984:Q2 to 2016:Q1, i.e., the period which includes all the regime changes in GI1 and can hence be considered as an unstable period. The relative MSFEs (RMSFEs), i.e., the ratio MSFE of Model 2 to Model 1 has been reported in Table 2 for forecast horizons of one-, two-, three-, and four-quarters-ahead. Understandably, if the RMSFE is less than one, then GD produces out-of-sample forecasting gains for GI1 over and above the other controls. As observed from the table, RMSFEs is less than one for horizons (b) of one-, two-, and three-quarters-ahead, and virtually the same for $b=4$. However, an important question is whether the forecasting gains produced by GD, i.e., Model 2 is statistically significant relative to Model 1. Since Model 2 nests Model 1, we use the *MSE-F* test statistic of McCracken (2007) to establish this point. Note that, the *MSE-F* statistic tests the null hypothesis that the restricted (Model 1) and unrestricted (Model 2) models have equal forecasting ability, against the one-sided alternative hypothesis that the MSFE for the unrestricted model is less than the MSFE for the restricted model forecasts. As depicted in Table 2, the *MSE-F* statistic was found to be significant at least at the 5% level of significance for $b = 1, 2$ and 3. Overall then, GD can produce significant forecasting gains for GI1 beyond the predictive contents of GRGDP, RHR, and RIR, especially for short- to medium-runs.

[INSERT TABLE 2]

In sum, GD has (time-varying) in (and out-of-sample) predictive content, particularly post the GFC, for GI1, with the sign of the effect being positive over the entire sample period.

IV. Conclusion

Theoretical and applied macroeconomic research have acknowledged household debt as a significant driver of business cycles. An interesting observation is that the observed credit booms from the data was mainly concentrated among the households outside the top of the income distribution. Consequently, in this study we explore the predictive power of the growth of household debt for growth in income (and consumption) inequality in the UK. We focus on the UK since income (and consumption) inequality data are available at a high frequency, i.e., on a quarterly basis for over 40 years (March, 1975 to March, 2016). Given that inequality is not only a problem in itself, but it also has negative economic, social, and health implications (Pierdzioch et al., 2019), we consider the usage of quarterly data to be of tremendous importance, since accurate prediction of inequality at a higher frequency should be more relevant to policymakers in designing appropriate policies to circumvent the wide-ranging negative impacts of uncertainty, compared to when forecasts are only available at the lower annual frequency.

Our findings point that higher growth rate in household debt correspond with subsequent increases in inequality, with growth in debt containing important information in predicting growth in income inequality for both in and out-of-sample periods. Interestingly, the strength of the in-sample predictive power is observed to have increased after the 2008 period. We believe these findings highlight an important role of household debt for inequality that was mostly ignored.

References

- Bai, J., and Perron, P. (2003). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, 18(1), 1-22.
- Barba, A., and Pivetti, M. (2009). Rising household debt: Its causes and macroeconomic implications—a long-period analysis. *Cambridge Journal of Economics*, 33(1), 113-137.
- Boivin, J., and Giannoni, M.P. (2006). Has monetary policy become more effective? *The Review of Economics and Statistics*, 88(3), 445-462.
- Campbell, J.Y., (2008) Viewpoint: estimating the equity premium. *Canadian Journal of Economics*, 41, 1–21.
- Clark, T.E., and McCracken, M.W. (2006). The predictive content of the output gap for inflation: resolving in-sample and out-of-sample evidence. *Journal of Money, Credit and Banking*, 38(5), 1127-1148.
- Cynamon, B. Z., and Fazzari, S.M. (2013). Inequality and household finance during the consumer age. *Levy Economics Institute of Bard College, Working Paper No. 752*.
- Del Negro, M., and Primiceri, G. (2015). Time Varying Structural Vector Autoregressions and Monetary Policy: A Corrigendum. *Review of Economic Studies*, 82(4), 1342-1345.
- Del Negro, M., and Schorfheide, F. (2013). DSGE Model-Based Forecasting. In: G. Elliott and A. Timmermann (eds.), *Handbook of Economic Forecasting, Volume 2, Part A*, p.57 – 140, Elsevier, New York.
- Dorling, D. (2015). Income inequality in the UK: Comparisons with five large Western European countries and the USA. *Applied Geography*, 61, 24-34.
- Iacoviello, M. (2008). Household debt and income inequality, 1963–2003. *Journal of Money, Credit and Banking*, 40(5), 929-965.
- McCracken, M.W. (2007) Asymptotics for out-of-sample tests of Granger causality. *Journal of Econometrics*, 140, 719–752.
- Mumtaz, H., and Theophilopoulou, A. (2017). The impact of monetary policy on inequality in the UK. An empirical analysis. *European Economic Review*, 98, 410-423.
- Pierdzioch, C., Gupta, R., Hassani, H., and Silva, E.S. (2019). Forecasting Changes of Economic Inequality: A Boosting Approach. *The Social Science Journal*. DOI: <https://doi.org/10.1016/j.soscij.2019.09.001>.
- Primiceri, G.E. (2005). Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, 72, 821-852.
- Rajan, R.G. (2010). *How hidden fractures still threaten the world economy*. Princeton, New Jersey.
- Rossi, B. (2005). Optimal tests for nested model selection with underlying parameter instability. *Econometric Theory*, 21(5), 962-990.
- Rossi, B. (2013). Advances in Forecasting under Model Instability. In: G. Elliott and A. Timmermann (eds.), *Handbook of Economic Forecasting, Volume 2B*, Elsevier, 1203-1324.

Saez, E. (2017). Income and wealth inequality: Evidence and policy implications. *Contemporary Economic Policy*, 35(1), 7-25.

Wang, Y., and Rossi, B. (Forthcoming). VAR-Based Granger-Causality Test in the Presence of Instabilities. *Stata Journal*.

Shi, S., Phillips, P.C.B., and Hurn, S. (2018). Change detection and the causal impact of the yield curve. *Journal of Time Series Analysis*, 39(6), 966-987.

Shi, S., Hurn, S., and Phillips, P.C.B. (Forthcoming). Causal Change Detection in Possibly Integrated Systems: Revisiting the Money-Income Relationship. *Journal of Financial Econometrics*.

Wood, J.D.G. (2019). Can household debt influence income inequality? Evidence from Britain: 1966–2016. *The British Journal of Politics and International Relations*. DOI: <https://doi.org/10.1177/1369148119888830>.

Table 1. Constant parameter and time-varying parameter Granger causality tests

| | $\chi^2(4)$ | <i>ExpW</i> | <i>MeanW</i> | <i>Nyblom</i> | <i>SupLR</i> |
|-----------------|-------------|-------------|--------------|---------------|--------------|
| Test Statistic | 13.1722 | 605.6220 | 1806.6731 | 123.5144 | 2483.3125 |
| <i>p</i> -value | 0.0105 | 0.0000 | 0.0000 | 0.0000 | 0.0000 |

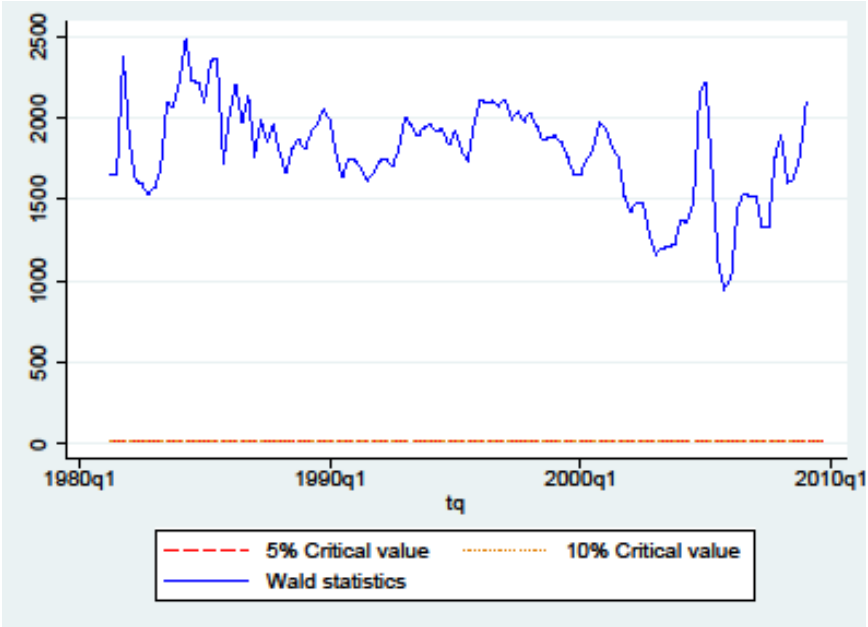
Note: Null hypothesis is the GD does not Granger cause GI1 in a constant or time-varying VAR(4), which includes GRGDP, RHR and RIR as additional variables. GI1 growth of Gini coefficient corresponding to income inequality; GD: growth in household debt; GRGDP: growth of real GDP; RHR: real housing returns, and; RIR: real interest rate.

Table 2. Out-of-sample forecasting of results of growth in income inequality (Gini coefficient)

| | Forecasting Horizon (<i>h</i>) | | | |
|--------------|----------------------------------|---------|---------|---------|
| | 1 | 2 | 3 | 4 |
| RMSFE | 0.9552 | 0.9534 | 0.9762 | 1.0006 |
| <i>MSE-F</i> | 5.9969* | 6.2068* | 3.0682# | -0.0760 |


Note: See Notes to Table 1; RMSFE is the ratio of MSFE of Model 2 to Model 1, where Model 1 includes GI1, GRGDP, RHR, and RIR, and Model 2 comprises of GI1, GRGDP, RHR, RIR, and GD; * and # corresponds to significance of the *MSE-F* test statistic at 1 and 5% levels respectively.

Figure 1. Time-varying Wald statistics with VAR(4) under AIC, testing whether GD Granger-causes GI1



Note: See Notes to Table 1; tq: corresponds to quarterly data; and the vertical axis measure the test statistic.


Figure 2. Time-varying Wald statistics with VAR(1) under SIC, testing whether GD Granger-causes GI1




Note: See Notes to Table 1 and Figure 1.

Figure 3: Results for recursive, rolling and recursive-rolling methods, testing whether GD Granger-causes GI1


(a) Recursive



(b) Rolling




(c) Recursive-Rolling



Note: See Notes to Table 1.

Figure 4. Response of GI1 to a one standard deviation GD shock in the TVP-VAR model



Note: See Notes to Table 1 and Figure 1.

APPENDIX:

Figure A1(a). Time Series Plot of Household Debt and Income Inequality (Gini Coefficient)






Figure A1(b). Time Series Plot of Household Debt and Consumption Inequality (Gini Coefficient)




Note: Period covered is 1975:Q1 to 2016Q1; Left vertical axis corresponds to household debt to GDP ratio, while right vertical axis measures the Gini coefficient.

Figure A1(c). Data Plots




Note: $GI_j, j=1,..6$, corresponds to the six measures of income and consumption inequality (Gini coefficient, the standard deviation, and the difference between the 90th and 10th percentile) respectively; GD: growth in household debt; GRGDP: growth of real GDP; RHR: real housing returns, and; RIR: real interest rate.

Figure A2. Time-varying Wald statistics with VAR(4) under AIC, testing whether GI1 Granger-causes GD



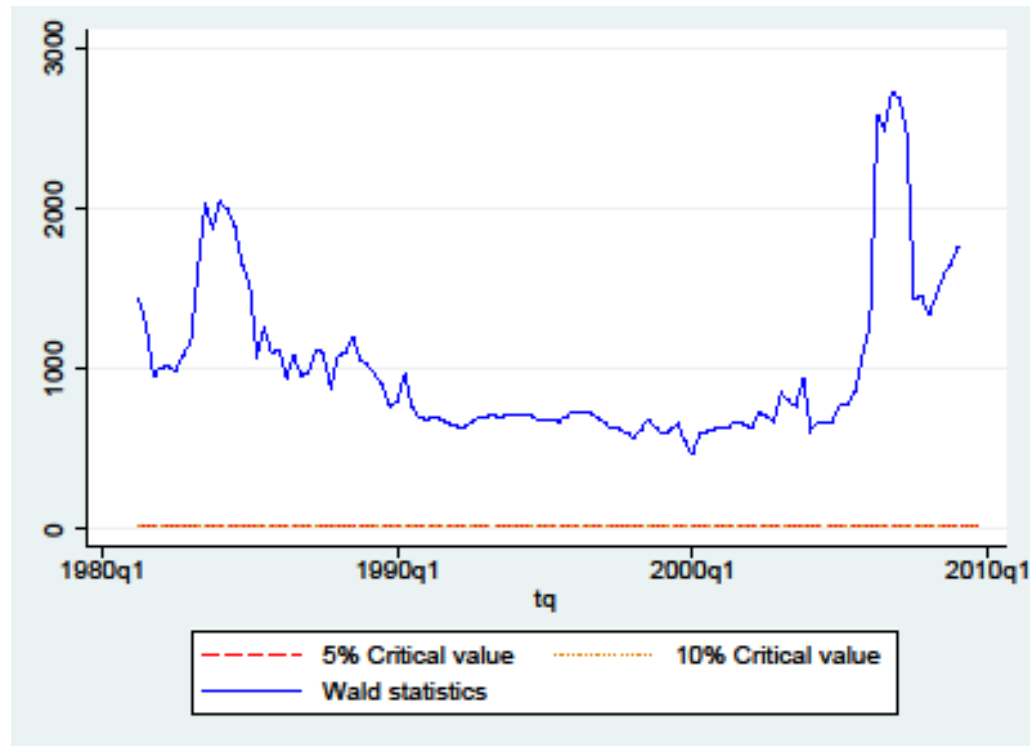
Note: See Notes to Figure A1(c); tq: corresponds to quarterly data; and the vertical axis measure the test statistic.

Figure A3. Time-varying Wald statistics with VAR(4) under AIC, testing whether GD Granger-causes GI2




Note: See Notes to Figure A1(c) and Figure A2.

Figure A4. Time-varying Wald statistics with VAR(4) under AIC, testing whether GD Granger-causes GI3



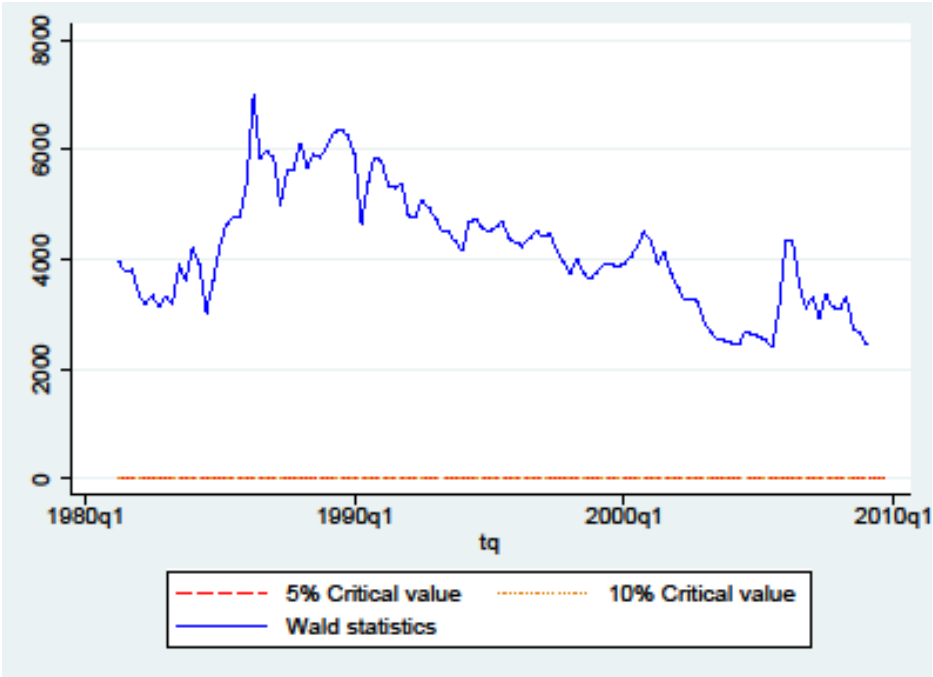
Note: See Notes to Figure A1(c) and Figure A2.

Figure A5. Time-varying Wald statistics with VAR(4) under AIC, testing whether GD Granger-causes GI4




Note: See Notes to Figure A1(c) and Figure A2.

Figure A6. Time-varying Wald statistics with VAR(4) under AIC, testing whether GD Granger-causes GI5



Note: See Notes to Figure A1(c) and Figure A2.

Figure A7. Time-varying Wald statistics with VAR(4) under AIC, testing whether GD Granger-causes GI6



Note: See Notes to Figure A1(c) and Figure A2.