Chaire en macroéconomie et prévisions

Predicting Financial Crises: Debt versus Debt Service Ratios

Steve Ambler[1] et Jeremy Kronick[2]

Cahier de recherche
Working paper
2020-09

Mai / May 2020

ESG UQÀM
École des sciences de la gestion
Université du Québec à Montréal

[1] Université du Québec à Montréal, C.D. Howe Institute
[2] C.D. Howe Institute

Chaire en macroéconomie et prévisions

La Chaire en macroéconomie et prévisions est fière de s'appuyer sur un partenariat avec les organisations suivantes:

Les opinions et analyses contenues dans les cahiers de recherche de la Chaire ne peuvent en aucun cas être attribuées aux partenaires ni à la Chaire elle-même et elles n'engagent que leurs auteurs.

Opinions and analyses contained in the Chair's working papers cannot be attributed to the Chair or its partners and are the sole responsibility of the authors.

Predicting Financial Crises: Debt versus Debt Service Ratios*

Jeremy Kronick
C.D. Howe Institute
jkronick@cdhowe.org

Steve Ambler
ESG UQAM, C.D. Howe Institute
ambler.steven@uqam.ca

March 2020

Abstract

Canada is often cited as having worryingly high credit-to-GDP and credit-to-disposable-income ratios, in spite of the fact that the assets and net worth of Canadian households have grown more quickly than their debt. We show that the level of debt servicing is a more reliable indicator of financial vulnerability than the level of debt itself. First, we construct a new financial vulnerabilities barometer and show that measures of debt servicing improve its ability to track periods of financial vulnerability, particularly in advance of recessions. Then, we show that the debt service ratio is a better predictor than the debt ratio of future declines in economic activity and banking crises. New borrowing, while supportive of economic growth in the short run, leads to an increase in debt servicing which contributes to slumps in economic activity.

*We thank participants at the annual meeting of the Société canadienne de science économique in Quebec City in May 2019 for helpful comments and discussions. Remaining errors are the responsibility of the authors.

1 Introduction

For the much of the last decade, Canadians have heard that their debt levels were unsustainable, and that their day of reckoning was fast approaching. According to data from the Office of the Superintendent of Bankruptcy (OSB), insolvencies by Canadian consumers were up 8.9 percent over a 12-month period ending in November 2019, compared to the 12-month period ending in November 2018.[1]

On the other hand, the data also show that Canadians' net worth has never been higher. Before the financial crisis, household net worth was equal to a little over 7 times disposable income. It dipped during the financial crisis, bottoming out at 6.25 times disposable income. However, since then, with low interest rates boosting housing prices, net worth has been on a mostly continuous upward trend, and has been between 8.5 and 9 times disposable income since the second quarter of 2016 (in the third quarter of 2019 net worth was 8.7 times disposable income). Moreover, while debt to GDP has increased from 75% in 2007 to 100% today, debt to net worth is almost identical at around 20%.[2] If borrowing is used to acquire capital assets or consumer durables or to start new businesses, debt levels or ratios of debt to income should not be a cause for worry.[3]

The OSB data do not distinguish between "cash flow insolvency" and "balance sheet insolvency." The second form of insolvency entails a negative net asset position. On the other hand, cash flow insolvency refers to the difficulty meeting payments on one's gross debt because of the difficulty selling one's assets. This type of insolvency is commonly referred to as illiquidity.

[1] See https://www.ic.gc.ca/eic/site/bsf-osb.nsf/eng/br04178.html.

[2] Data from Statistics Canada 38-10-0235-01. See Ambler, Kronick and Omran (2019) for more.

[3] Congdon (2017) notes that the OECD and the BIS have been ringing similar alarm bells about debt levels around the world while net assets have been increasing. "The historical experience is that assets increase by amounts so much larger than debt that the net-wealth-to-income ratio and the debt-to-income ratio can and usually do rise together!"

Unfortunately, the only distinction in the OSB data is between bankruptcies and consumer proposals, an alternative to bankruptcy where the lender and borrower strike a deal under which the latter pays back a portion of their debt over the ensuing five years (and, critically, does not have to sell their house). In the 12-month period ending November 2019, consumer proposals increased 17.4 percent while bankruptcies decreased 1.9 percent. With stable debt to net worth, and with two-thirds of Canadian households owning homes, it could very well be that the bulk of these consumer proposals are a result of illiquidity concerns. In this case, the debt service burden will be an important predictor of insolvencies.

In this paper, we construct a financial vulnerability barometer for Canada in which debt servicing plays a key role. We find that debt servicing improves the barometer's ability to predict periods of financial vulnerability, especially preceding recessions. We then go on to investigate the statistical significance of debt servicing for predicting financial vulnerability and economic growth. We develop a simple model of the dynamics of debt and debt servicing based on the work of Drehmann, Juselius and Korinek (2017). We take this model to Canadian quarterly data. The model allows us to simulate the dynamics of the debt service ratio in response to an innovation in new borrowing using the local projection method of Jorda (2005).

Our results show that new borrowing is highly persistent, and that debt servicing peaks with a lag after a positive innovation to borrowing. We find that increases in new borrowing lead initially to a slight (and barely significant) increase in economic activity and a significant reduction in the probability of financial crisis. The debt service ratio has a significantly negative impact on economic activity and a significantly positive impact on the probability of financial crisis.

We also run a horse race to test the importance of new borrowing and debt servicing compared with more traditional credit variables. We run regressions for the

1990:Q1–2019:Q3 period that are designed to forecast one-period-ahead real GDP growth and the probability of financial crises (also one-period ahead) to see which variables are more important for predicting growth and financial crises. New borrowing and debt service are significant explanatory variables for both GDP growth and financial crises, with signs confirming our earlier results and coefficients that are quite stable across specifications. The same cannot be said for more traditional credit measures. Moreover, new borrowing and debt servicing explain the highest amount of variation of the dependent variables as measured by the R^2.

What this implies is that debt servicing is a much better predictor of the health of an economy than is debt/GDP or other ratio measures that conflate stock (e.g. debt) and flow (e.g. GDP) variables. Given how flat the debt service ratio has been over the last 25 years, despite growth in the debt/GDP ratio, these results might explain why we have not seen the type of market correction that has been repeatedly predicted for the Canadian housing market.

In the following section, we describe the construction of our financial vulnerability barometer. In the third section, we look at the accounting relationship between new borrowing and debt service and some descriptive statistics on the dynamics of these two variables. In the fourth section, we estimate a formal statistical model of the dynamic relationship between new borrowing and debt servicing and look at the impulse responses of both new borrowing and debt servicing to an innovation in new borrowing. In the fifth section, we estimate the impact of new borrowing and debt servicing on economic activity and the probability of financial crisis, while in section six we look at the out-of-sample predictive power of new borrowing and debt servicing for GDP growth and financial vulnerability. The seventh section discusses some robustness tests, while the final section draws policy implications and conclusions. A technical appendix discusses details of the construction of the financial vulnerability index.

2 A New Financial Vulnerability Barometer

Figure 1 below shows that the debt-service ratio of Canadian households has been quite stable over our entire sample period, except for a mild recent increase, and, critically, in the run-up to the financial crisis. The stability largely comes from increases in principal, offset by decreases in interest rates. The figure indicates that increases in the debt-service ratio may help to predict financial stress associated with recessions, in particular the financial crisis and Great Recession of 2007–08.

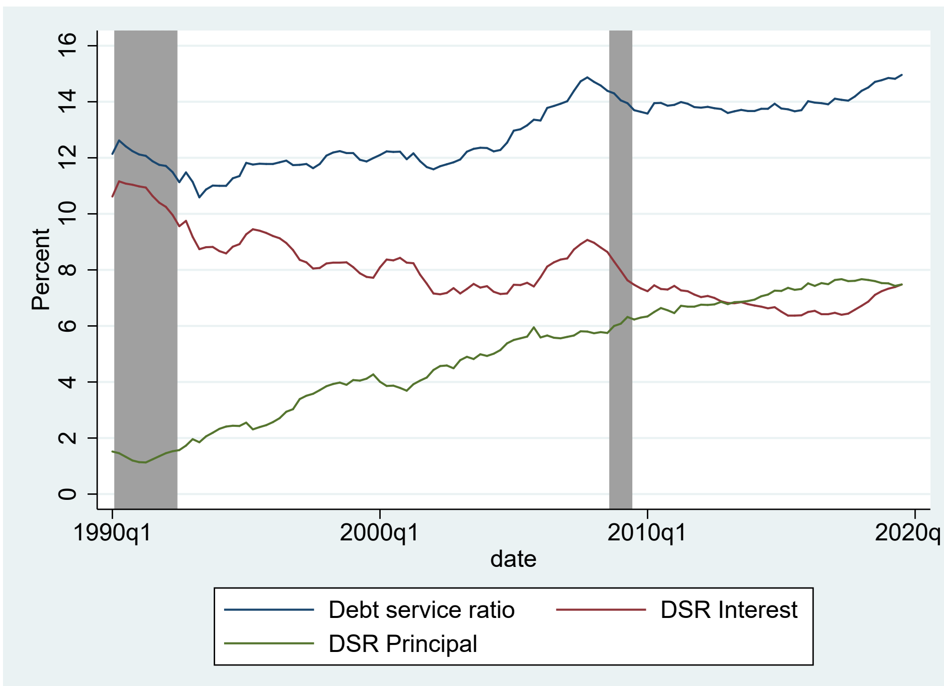

Figure 1: Debt service ratio

While debt servicing data are not available before 1990, it is possible to generate an estimated version of debt servicing going further back in time.[4] Figure 2 backdates debt servicing to 1980. The estimated debt-service ratio peaked in advance of the early 1980s recession, and increased sharply prior to the early 1990s recession. This reinforces the idea

[4]See Drehmann and Juselius (2012) for details.

that changes in this variable may help predict financial stress.[5]

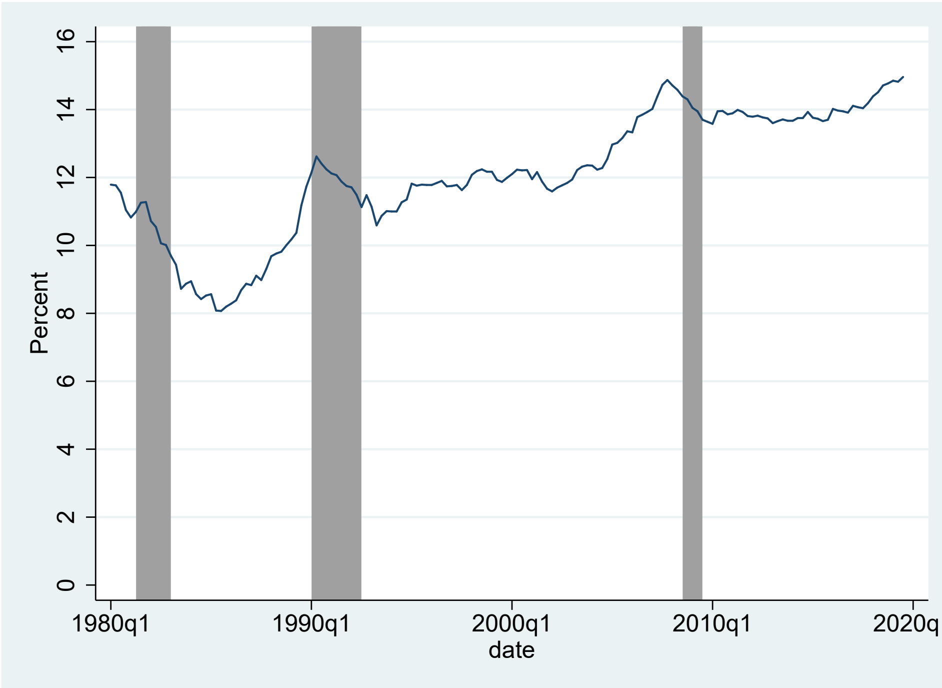

Figure 2: Debt service ratio - estimates pre-1990

The Bank of Canada[6] has developed a financial stability vulnerabilities barometer which aggregates together indicators from the banking, corporate, household, and housing sectors. The choice of indicator variables is based on the methodology of "receiver operating curves" (ROC), which scores indicators on their ability to predict periods of financial stress within 24 months. We explain the principles behind this methodology in the technical appendix below.[7] Only indicators that scored greater than 0.5 are retained: an indicator with a lower score does worse than a coin flip for predicting financial stress.

In this section, we compare the evolution of the Bank's financial stability barometer with one that incorporates debt servicing.

[5] As a result of the backdating methodology, estimated data pre-1990 were scaled down to align with actual post-1990 data. The correlation between actual post-1990 data and estimated post-1990 data if we were to extend the backdating methodology forward is very high (above 0.9).

[6] See Pasricha et al. (2013) and, more recently, Duprey and Roberts (2017), which is our focus.

[7] See also Fawcett (2006), Flach (2010), Swets et al. (2000) and van Erkel and Pattynama (1998) for more technical details.

The Duprey and Roberts (2017) vulnerabilities barometer is constructed using the following equation:

$$VB_t = \sum_{s=1}^{S} \max \left(\sum_{i=1}^{I_s} \max \left(\frac{v_{s,i,t} - \tau_{s,i}}{\sigma_{s,i}}; -1 \right) * w_{s,i}; 0 \right). \tag{1}$$

Here, v is the value of the indicator i in sector s at time t. A financial stress episode occurs when a financial stress index is two standard deviations above its trailing ten-year moving average. Thresholds associated with stress periods are calculated by minimizing the "impurity measure", which measures the extent to which signals issued by indicators that are above their threshold mis-classify stress periods. The optimal threshold also minimizes the noise-to signal ratio of the indicator: see the technical appendix for further details. The weights, $w_{s,i}$, give the relative informational content of a given indicator relative to other indicators in that sector. Duprey and Roberts use the area under the receiver operating curve (AUROC),[8] defined in the appendix, to calculate these weights, using only those that are above 0.5. Weight $w_{s,i}$ (for sector s) is defined as

$$w_{s,i} = \frac{\text{AUROC}_{s,i} - 0.5}{\sum_i \left(\text{AUROC}_{s,i} - 0.5\right)}$$

Because financial stress episodes are rare in individual countries, we follow Drehmann and Juselius (2014) and look at international crises periods across a panel of countries including Canada over the 1990:Q1–2019:Q3 period. The exact dates can be found in Drehmann and Juselius (2014), to which we add the oil price shock and economic slowdown in 2015:Q1. These financial periods correspond to the following events:

1. The recession of the early 1990s

2. The Mexican crisis (1994–1995)

[8]This gives a measure of how well the indicator variable does compared to a coin flip across different possible threshold values.

3. The Asian crisis (1997–1998)
4. The Russian debt default (1998)
5. The Long-Term Capital Management collapse (1998)
6. The sub-prime crisis, financial crisis and Great Recession (2008–2009)
7. The aftermath of the oil price collapse (Early 2015)

We simplify the Duprey-Roberts measure in a variety of ways, and focus on the additional value-added of incorporating debt servicing. Rather than using a panel of countries, we calculate our own made-in-Canada thresholds.

We do not use all variables across sectors from Duprey and Roberts. Instead, we begin by focusing on the two highest-ranked AUROCs in their four categories, excluding those that conflate stock and flow variables, e.g. pure debt/GDP ratios. We also focus on deviations from a one-sided Hodrick-Prescott filter, as a regulator analyzing a vulnerabilities barometer at a point in time would only have backward-looking information (see Drehmann and Juselius (2014) for more). We are left with the following variables:

- the one-sided deviation of the ratio of household debt to disposable income;
- the one-sided deviation of the ratio of household debt to GDP;
- the ratio of housing price to rent;
- the one-sided deviation of the ratio of housing price to rent;
- the year-over-year growth rate of the ratio of non-financial corporate debt to GDP;
- the one-sided deviation of the ratio of non-financial corporate debt to GDP;
- the year-over-year growth rate of the ratio of financial institution debt to GDP;

Table 1: Thresholds for Variables of Interest

Variable	Min-Max	Threshold	AUROC
HH Debt/DI Dev from Trend	-8.79 - 13.04	0.01	0.71
HH Debt/GDP Dev from Trend	-3.97 - 14.35	6.83	0.66
HP/Rent Dev from Trend	-2.80 - 16.30	8.20	0.73
Non-Fin Corp Debt/GDP YoY Gr	-8.77 - 13.87	0.23	0.64
Non-Fin Corp Debt/GDP Dev from Trend	-12.23 - 17.03	7.77	0.60
FI Debt/GDP Dev from Trend	-7.85 - 15.27	4.35	0.61
FI Debt/GDP YoY Gr	-20.95 - 39.00	10.40	0.60
DSR Dev from Trend	-0.60 - 1.57	0.80	0.62

- the one-sided deviation of the ratio of financial institution debt to GDP.

The AUROCs in Duprey and Roberts (2017) come from estimation on a panel of countries. We test these variables using only Canadian data as well to see if the AUROCs remain greater than 0.5. One variable does not: the ratio of housing price to rent. We, therefore, drop it from the sample.

Drehmann and Juselius (2014) find a very high AUROC for the debt service ratio across a panel of countries (> 0.8). However, when we test on Canadian data alone, the AUROC falls below 0.5. However, the AUROC for the deviation of the debt service ratio from a one-sided Hodrick-Prescott filter is well above 0.5 (0.62). We therefore add this variable to our list. Our final list of eight variables can be found in Table 1 below, with their associated thresholds and the AUROCs from Duprey and Roberts, and in the case of the DSR deviation, from our own estimation.[9]

The resulting vulnerabilities barometer is shown in Figure 3 below. The grey bars are Canadian recessions and the yellow bars represent international financial stress episodes.

[9] We deem that the Duprey and Roberts (2017) deviation from a 10-year moving average is sufficiently similar to our one-sided Hodrick-Prescott filter deviation to be able to use their calculated AUROCs.

Figure 3: Vulnerabilities Barometer

The barometer does a good job of tracking both major recessions and other periods of elevated financial stress. Debt servicing also appears to have significant value added, especially in advance of recessionary periods. Also of interest (and in contrast to the Duprey and Roberts measure) is that it indicates that there is no point post-crisis when we reached the level seen before the financial crisis. As of 2019:Q3 risks appear low.[10] The Duprey and Roberts measure does well in capturing all crisis periods, but it also results in quite a few false alarms. Our measure appears to mitigate some of these false positives.

Data exist going back further than 1990 for household debt to disposable income, household debt to GDP, and the house price to rent ratio. As in Drehmann and Juselius (2012), this allows us to calculate a debt service ratio a vulnerabilities barometer extending back to 1980. Two financial stress periods are added to this larger data set: the recession

[10]The recent COVID-19 event would not be captured by our dataset. Moreover, we would not expect any kind of financial vulnerabilities barometer to pick up a black swan event such as this global pandemic.

in Canada in the early 1980s and the 1987 stock market crash, the latter consistent with Drehmann and Juselius (2014). The extended barometer is shown in Figure 4 below. Once again, it does a good job of picking up these stress periods, and the debt service ratio provides significant value added in this earlier period.

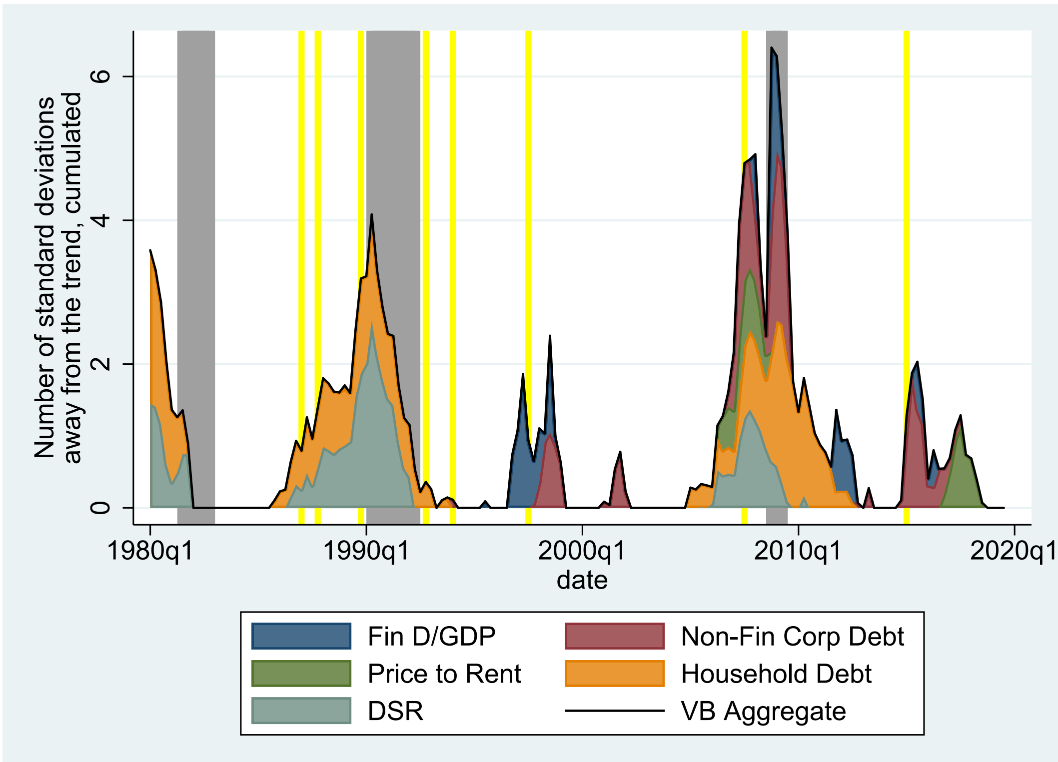

Figure 4: Extended Vulnerabilities Barometer

Our vulnerability barometer establishes the *prima facie* usefulness of the debt service ratio as a predictive tool. However, formal tests of the statistical significance of debt servicing to forecast economic growth and fragility are difficult to implement with our constructed measures. To do this we need to use more formal econometric techniques, to which we now turn.

3 New Borrowing and Debt Servicing - An Accounting Framework

We use a simple accounting framework based on Drehmann, Juselius and Korinek (2017) to document the persistence of new borrowing and the dynamic relationship between new borrowing and debt servicing.

New borrowing is measured as follows. We begin with seasonally-adjusted data from Statistics Canada on nominal credit. We then use seasonally-adjusted data on principal payments and divide these by nominal credit to get the amortization rate. If D_t is the stock of debt at the beginning of t and δ_{t-1} is the amortization rate on the stock of debt in $t-1$ then debt must follow the following simple dynamic relationship:

$$D_t = (1 - \delta_{t-1})\, D_{t-1} + B_{t-1} \tag{2}$$

where B_{t-1} is new borrowing during $t-1$. Total debt servicing including interest and payments on principal is given by

$$S_t = (\delta_t + r_t)\, D_t \tag{3}$$

where r_t is the average (real) interest rate on outstanding debt. Net cash flow, N_t, from lenders to borrowers is given by new borrowing less total payments to service the debt:

$$N_t = B_t - S_t. \tag{4}$$

By successive substitution, equations (2) and (3) imply the following dynamic relationship between new borrowing and debt service:

$$S_t = (\delta_t + r_t) \sum_{i=1}^{t} \prod_{j=1}^{i-1} (1 - \delta_{t-j})\, B_{t-i} \tag{5}$$

We scale new borrowing by seasonally-adjusted GDP and take logs. Our debt service variable is the total debt service ratio, also measured in logs.

If new borrowing is serially correlated, it can be shown that debt servicing must continue to increase for some time after peaks in new borrowing.[11] Since debt service is a function of the stock of debt, the stock of debt and debt service are still increasing when new borrowing peaks. After the peak, a lower amortization rate δ delays the time when debt service equals (declining) new borrowing.

Quarterly Canadian data exhibit precisely these broad patterns over time. The statistical properties of new borrowing and the debt service ratio are illustrated in Figures 5 and 6 below. Figure 5 clearly shows that new borrowing is highly serially correlated.[12]

[11] See Appendix A in Drehmann, Juselius and Korinek (2017) for a formal proof.

[12] The first-order autocorrelation coefficient of new borrowing is higher than that obtained by Drehmann, Juselius and Korinek (2017) for a cross section of countries. This is to be expected since they use annual data while we use a quarterly data set.

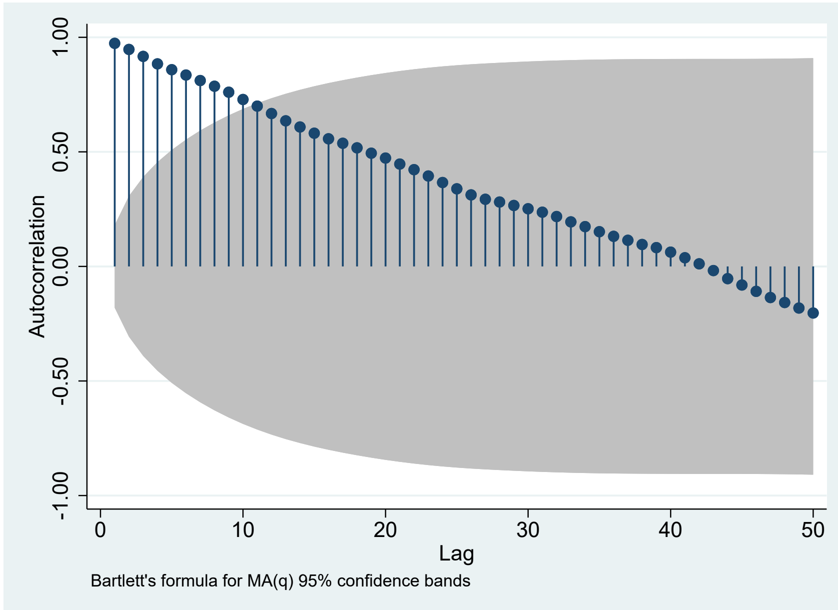

Figure 5: Autocorrelation Function of New Borrowing

The debt service ratio then follows new borrowing with a lag, as shown by the cross-correlogram in Figure 6 below. The peak correlation between debt servicing and new borrowing occurs at a five-quarter lag.

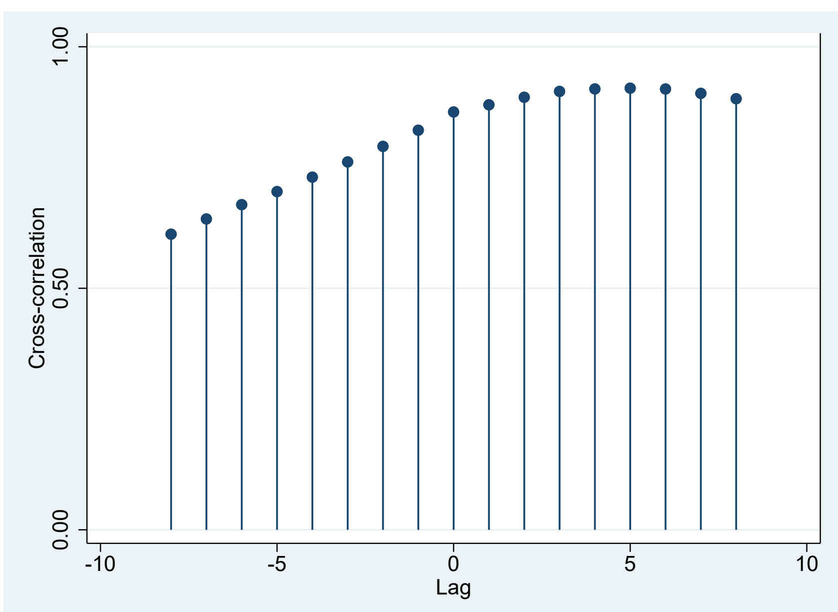

Figure 6: Cross-correlation between new borrowing and debt service

Figure 7 is another way of looking at the dynamic relationship between new borrowing and debt service. It shows the average evolution of the two variables around peaks in new borrowing. A peak in new borrowing is defined as a local maximum within a three-year window. The figure shows quite clearly that peaks in the debt service ratio follow peaks in new borrowing.

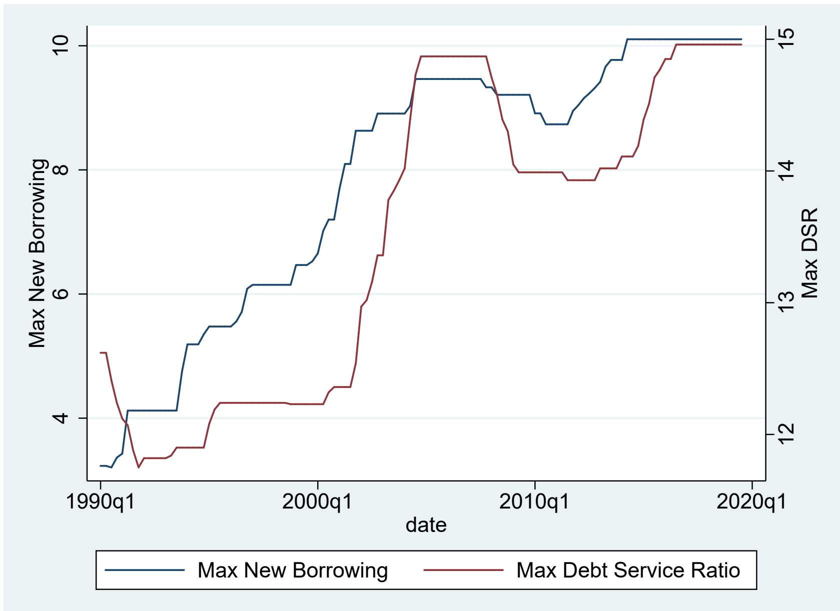

Figure 7: Evolution of new borrowing and debt service around peaks in the former

4 The Impact of New Borrowing on Itself and the Debt Servicing Ratio

In order to study more formally the relationship between new borrowing and debt servicing, we follow Drehmann, Juselius, and Korinek (2017), and estimate the following equations.

$$b_{t+h} = \beta_0 + \beta_{bb}^h b_t + \beta_{bs}^h s_t + \text{controls} + \epsilon_{b,t+h}^h \tag{6}$$

$$s_{t+h} = \beta_0 + \beta^h_{sb} b_t + \beta^h_{ss} s_t + \text{controls} + \epsilon^h_{s,t+h} \tag{7}$$

The control variables we use are described in detail in the data appendix, and include a real short rate, a lending spread, a change in the average lending rate on household debt, real property price growth, and real GDP growth. We simulate the response of new borrowing and debt service to an innovation in net borrowing, using the local projection method of Jorda (2005). The h successive β^h_{bb} and β^h_{sb} coefficients give the impulse response of future new borrowing and future debt service, respectively, to a unit increase in new borrowing at time t over h successive years.

Figure 8 below shows the simulated responses of new borrowing and debt service to a positive unit innovation in new borrowing. It shows clearly that a positive shock to new borrowing persists over time and that debt servicing responds positively but with a lag to a positive innovation in new borrowing.[13] The dark grey shaded areas indicate 90% confidence bands. The light grey shaded areas are 95% confidence bands.

[13] These impulse responses are for regressions where no lags of new borrowing or debt servicing are included in the estimated equations. The results with one lag of each variable included as regressors are similar but less smooth. We also calculated the impulse responses of new borrowing and debt servicing to an innovation in the former using an estimated VAR and a Cholesky decomposition that allows debt servicing to respond to new borrowing only with a lag. Once again the results were qualitatively very similar. We present these results in the Robustness section below.

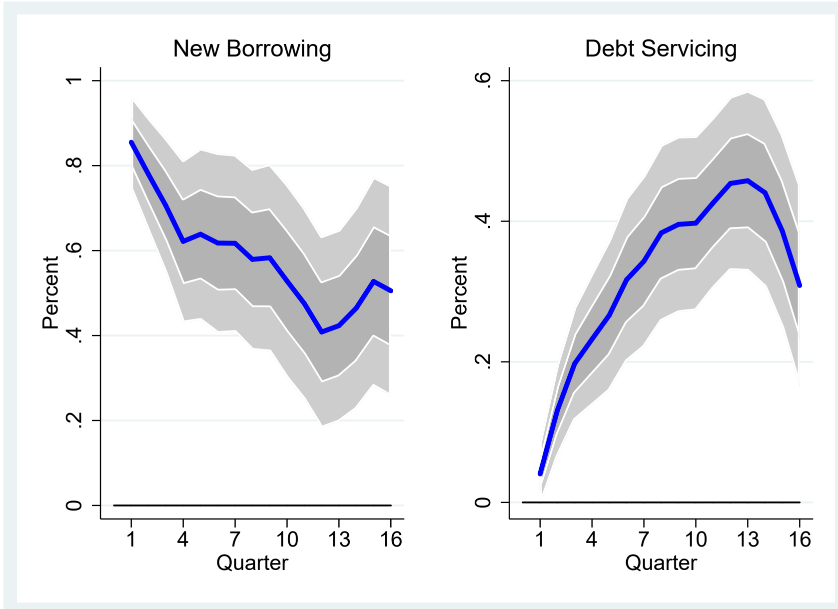

Figure 8: Impulse response of new borrowing and debt service to a unit increase in new HH borrowing

5 The Impact of New Borrowing and Debt Service on Growth and Financial Stability

In this section, we examine the dynamic relationship between new borrowing and debt servicing on the one hand and GDP growth and financial stability on the other hand.

We look first at real growth. Again using the methodology of Jorda (2005), and following Drehmann, Juselius, and Korinek (2017), we estimate local projections of GDP growth following a one unit shock to both new borrowing and debt servicing in order to assess the dynamic impact of these two variables on real growth. We estimate equations of the form

$$\Delta y_{t+h} = \beta_0 + \beta^h_{yb} b_t + \beta^h_{ys} s_t + \text{controls} + \epsilon^h_{y,t+h} \tag{8}$$

The control variables remain the same.

The estimates of β_{yb}^h and β_{ys}^h for successive values of h trace out the impulse response of GDP growth given unit increases in new borrowing and future debt service, respectively.

Figure 9 shows the results. The figure indicates that new borrowing has an initial positive effect on GDP growth, but the effect turns negative late in year two. Debt service, on the other hand, has an immediate negative impact on GDP growth. What this says is that while new borrowing may have an initial positive effect on GDP growth, it is short-lived, and the subsequent increase on debt servicing leads to a decrease in GDP growth.

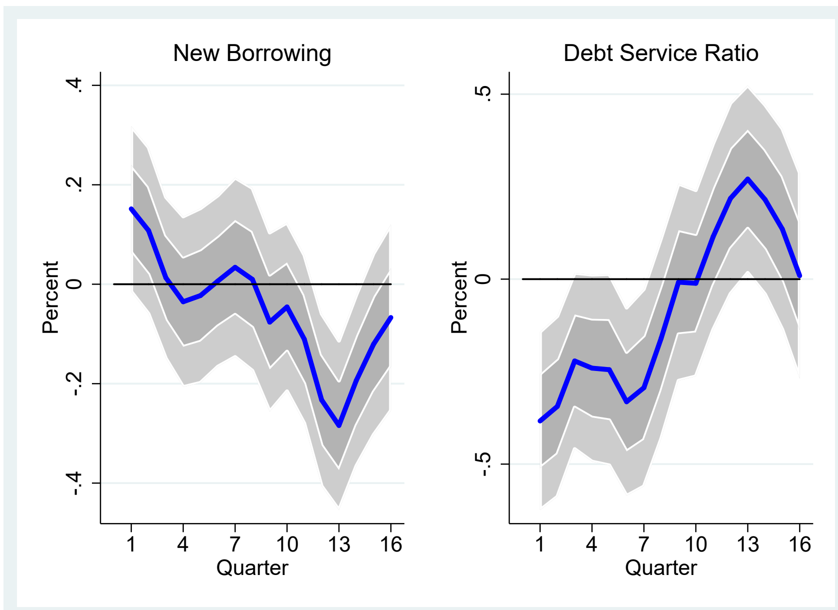

Figure 9: Impulse response of GDP growth after a unit increase in new HH borrowing or debt servicing

We can also estimate the dynamic response of the probability of financial crisis in response to new borrowing and the debt service ratio using Jorda's local projections method:

$$c_{t+h} = \beta_0 + \beta_{yb}^h b_t + \beta_{ys}^h s_t + \text{controls} + \epsilon_{y,t+h}^h, \tag{9}$$

where c_{t+h} is the value of our crisis indicator variable. This variable receives a 1 during recession quarters and a zero otherwise. For our data sample, there are two recessionary periods, 1990:Q2–1992:Q2 and 2008:Q4–2009:Q2.[14]

Figure 10 below shows the results of simulations based on these regressions. The response of the crisis indicator variable is similar to the response of real growth (with a reversal of sign). New borrowing initially has a negative impact on the probability of financial crisis, with a change in sign in the middle of year 2. The debt service ratio has an immediate and pronounced positive impact on the probability of crisis. The conclusion is, therefore, similar to the case of GDP growth: while new borrowing is initially good for the economy, it eventually leads to an increase in debt servicing, which leads to instability and increased risk of a financial crisis.[15]

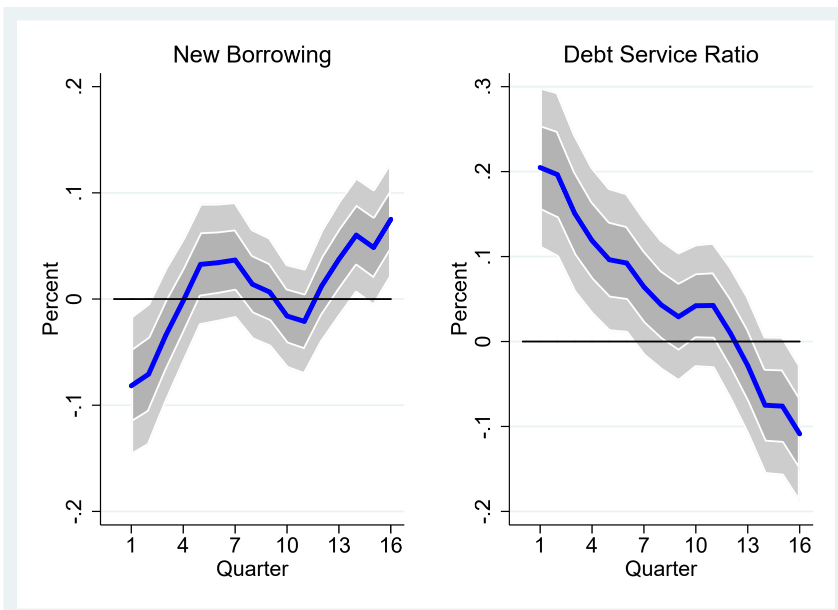

Figure 10: Impact of new borrowing and debt service on the probability of crises

[14]Drehmann, Juselius and Korinek (2017) for their panel use a series of crises across all countries, even if one individual country isn't affected by a particular crisis. Since we do not have a panel, we restrict our crises to Canadian recessionary periods, as defined by the C.D. Howe Institute's Business Cycle Council. The results are similar if we add additional crisis periods. Results available upon request.

[15]Our baseline local projections estimation is run using OLS. The results hold when we run a local projection using logit as well.

6 Out-of-sample predictive power

What happens if we run a horse race between new borrowing and debt servicing on the one hand, and more traditional measures of debt, such as credit-to-GDP, on the other? We run regressions for the 1990:Q1–2019:Q3 period that are designed to forecast one-period-ahead real GDP growth and the probability of financial crises (also one-period ahead) to see which variables are more important for predicting growth and financial crises, as defined above. Tables 2 and 3 below give an assessment of the relative predictive power for one-quarter-ahead GDP growth and one-quarter-ahead probability of a financial crisis. Table 3 uses ordinary least squares with the financial crisis variable as the dependent variable. A logit regression produces qualitatively similar results.[16]

Table 2: The effects of different credit measures on output growth

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
New Borr	0.117*					0.259***	0.174**	0.214**	0.072
	(0.062)					(0.069)	(0.076)	(0.086)	(0.058)
DSR	-0.295***					-0.458***	-0.306***	-0.415***	-0.234**
	(0.096)					(0.104)	(0.098)	(0.125)	(0.089)
Credit Gr		-0.178**				-0.275***			
		(0.080)				(0.075)			
Cr/GDP Gr			-0.010				-0.032		
			(0.026)				(0.035)		
Gap HP1				-0.010				-0.074	
				(0.038)				(0.049)	
Gap HP2					0.097***				0.085**
					(0.033)				(0.033)
Observations	117	117	114	117	117	117	114	117	117
R^2	0.42	0.42	0.32	0.37	0.42	0.51	0.37	0.44	0.45

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
Estimated with baseline controls as in local projections figures (no lags).
DSR = debt service ratio, Credit Gr = quarterly credit growth, Cr/GDP Gr = year over year growth in credit-to-GDP ratio, Gap HP1 = one-sided Hodrick-Prescott filter deviation of credit-to-GDP from trend, Gap HP2 = two-sided Hodrick-Prescott filter deviation of credit-to-GDP from trend.
Source: Authors' calculations.

[16]The details are available on request from the authors.

Table 2 shows that new borrowing and debt service are significant explanatory variables for GDP growth. As expected, new borrowing has a positive effect on future GDP growth while the debt service ratio has a strongly negative effect. Both variables are significant both on their own (column 1), and when more traditional variables are added to the regressions (columns 6-9). Moreover, the coefficients are fairly stable across specifications. These results confirm the results in the cross-country study by Drehmann, Juselius and Korinek (2017).[17] The results also remain true when we extend the forecast period of the dependent variable from one-quarter-ahead real GDP growth to one-year-ahead real GDP growth.[18]

Credit growth also appears to be significant when included on its own, and when added to regressions that include new borrowing and the debt-service ratio. In fact, the latter has the highest percentage of variance explained (as measured by the R^2) of any regression we run. The difference between credit growth and new borrowing can best be summarized with a quick example. Imagine the stock of credit this period is $100, as it was last period. Credit growth would be zero. However, if principal payments on last period's stock were $20, and new borrowers took out $20, new borrowing would equal $20.

Year-over-year growth in credit-to-GDP, on the other hand, is insignificant when added on its own as an explanatory variable, and when added to regressions that include new borrowing and the debt-service ratio.

We also look at two different credit gap measures. Specifically, using both a one-sided and two-sided gap, where we calculate the deviation of credit-to-GDP from its trend (Gap HP1 and Gap HP2, using Hodrick-Prescott filter). The difference between a one-sided and two-sided gap is that in the former case we only use backward-looking data, whereas in the

[17] They find that the impact of debt service remains stable across specifications and when different measures of debt are added. They find the new borrowing variable to be less stable, with multicollinearity affecting its estimated coefficient when different measures of debt are included.

[18] At two years ahead, the results remain true for debt servicing, while new borrowing, despite having the same sign, loses its significance.

latter we use both backward and forward-looking data. In the case of the one-sided gap variable, it is insignificant on its own, and when added to a regression including new borrowing and debt servicing. In the case of the two-sided gap, the coefficients are significant and positive in both, though this measure is unlikely to be as useful as new borrowing and debt servicing in forecasting, since a regulator only has backward-looking information to use in assessing vulnerabilities (see Drehmann and Juselius 2014 for more details on this point).

Table 3: The effects of different credit measures on crisis probability (OLS)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
New Borr	-0.070**					-0.105***	-0.122***	-0.119***	-0.070**
	(0.030)					(0.036)	(0.038)	(0.044)	(0.030)
DSR	0.165***					0.205***	0.187***	0.225***	0.166***
	(0.050)					(0.056)	(0.053)	(0.065)	(0.048)
Credit Gr		0.030				0.069**			
		(0.029)				(0.029)			
Cr/GDP Gr			0.008				0.026*		
			(0.010)				(0.014)		
Gap HP1				0.002				0.038**	
				(0.013)				(0.019)	
Gap HP2					-0.009				0.001
					(0.015)				(0.014)
Observations	117	117	114	117	117	117	114	117	117
R^2	0.55	0.49	0.39	0.48	0.48	0.58	0.49	0.57	0.55

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
Estimated with baseline controls as in local projections figures (no lags).
DSR = debt service ratio, Credit Gr = quarterly credit growth, Cr/GDP Gr = year over year growth in credit-to-GDP ratio, Gap HP1 = one-sided Hodrick-Prescott filter deviation of credit-to-GDP from trend, Gap HP2 = two-sided Hodrick-Prescott filter deviation of credit-to-GDP from trend.
Source: Authors' calculations.

Table 3 shows that the results are qualitatively similar to the GDP growth prediction equation, but with a sign reversal. Both new borrowing and debt servicing are highly significant, and their numerical values are stable across all specifications. New borrowing significantly reduces the probability of financial crisis[19] while the debt service ratio has a

[19]Drehmann, Juselius and Korinek (2017) find a significantly positive effect, though their finding is not

strongly positive impact on the probability of financial crisis. Traditional measures when looked at without controlling for new borrowing and debt service have no significance. And, while their significance improves when combined with new borrowing and debt service, they have only a mild effect on the percentage of variance explained (as measured by the equations' R^2) when compared to the regression with just new borrowing and debt service as explanatory variables.

7 Robustness

We performed a series of robustness checks. These checks all confirm our original findings. For the sake of brevity we present the result of three checks here.

The first check involved the substitution of the local projections methodology for a structural vector autoregression (SVAR), with a Cholesky decomposition. We use the same variables used in the local projections methodology, with one lag in the SVAR. Our vector of variables, y_t is ordered as follows: $y_t = [r_t, i_t, \Delta l_t, \Delta p_t, b_t, s_t, g_t]$, where, in addition to our new borrowing and debt servicing variables, b_t and s_t, we have our control variables, r_t, the real short rate, i_t, the lending spread, Δl_t, the change in the average lending rate on household debt, Δp_t, real property price growth, and g_t, real GDP growth (full definitions for variables can be found in the Data Appendix). For the estimation involving recessions we insert the crisis dummy (as described above) in the SVAR, in between debt servicing and real GDP growth. Critically, we order debt servicing after new borrowing, in all cases, to allow the former to respond with a lag to the latter.

Figures 11–13 replicate Figures 8–10, and the results are qualitatively the same.[20] A shock

robust to their different specifications. In fact, they find collinearity between new borrowing and the other measures, wiping away the significance of the new borrowing variable on financial crises

[20]Note that the magnitudes of the coefficients will be lower in these figures than in Figures 8–10. Whereas in the local projections methodology, the size of the shock is 100 basis points, it is based on a one unit

to new borrowing causes a persistent increase in new borrowing. The effect of this new borrowing shock on debt servicing is also positive, though the increase occurs with a lag (Figure 11).

Moreover, a shock to new borrowing causes an initially positive increase in real GDP growth, though turns negative in the second year. An increase in debt servicing has a negative effect on real GDP growth (Figure 12).

Lastly, an increase in new borrowing reduces the probability of a crisis, but the effect becomes (insignificantly) positive after 10 quarters. By contrast, and consistent with what we saw before, an increase in the debt servicing ratio causes an increase in the probability of a crisis (Figure 13).

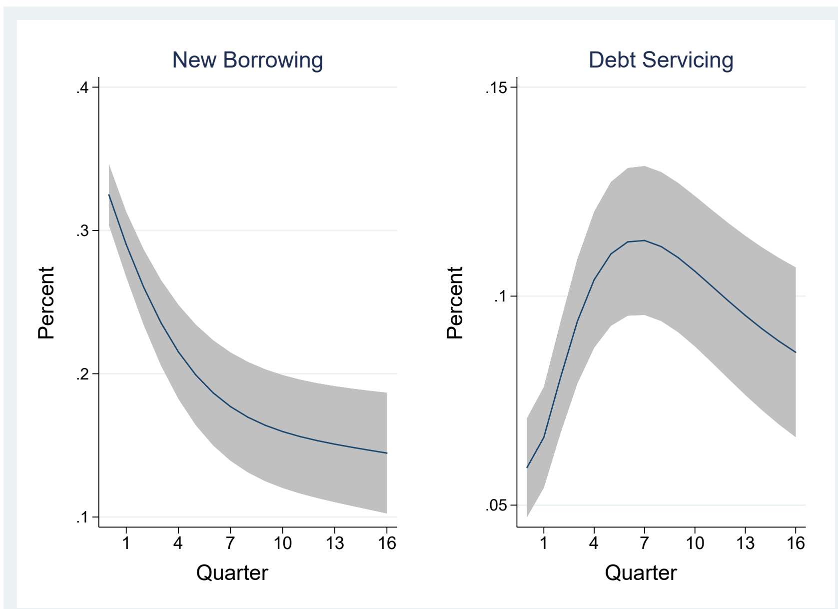

Figure 11: Impulse response of new borrowing and debt service to a unit increase in new HH borrowing

standard deviation in the case of the SVAR, which in this case is lower.

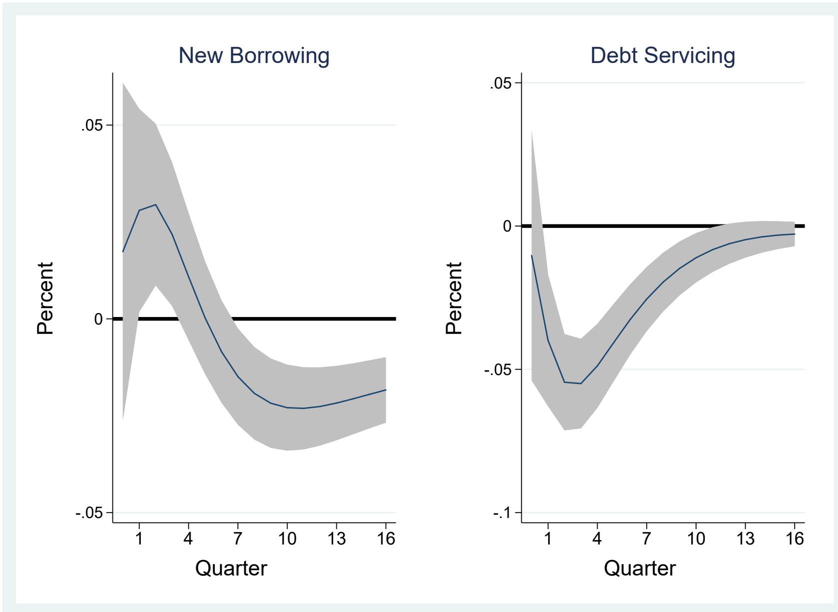

Figure 12: Impulse response of GDP growth after a unit increase in new HH borrowing or debt servicing

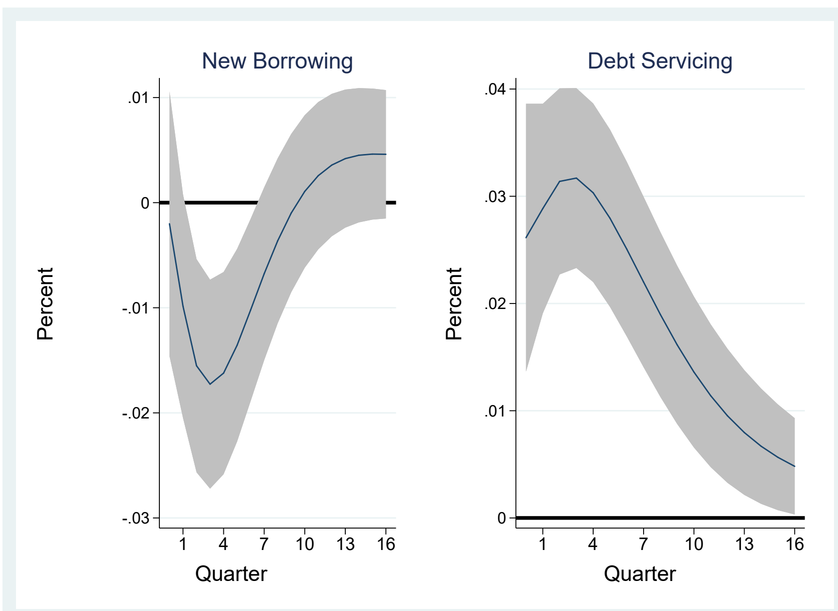

Figure 13: Impact of new borrowing and debt service on the probability of crises

As a second robustness check, we include more than just the two Canadian recessions as crisis periods, using the international crises periods from Drehmann and Juselius (2013), as described above. The first robustness check involves replicating Figure 10 with the new crisis variable.

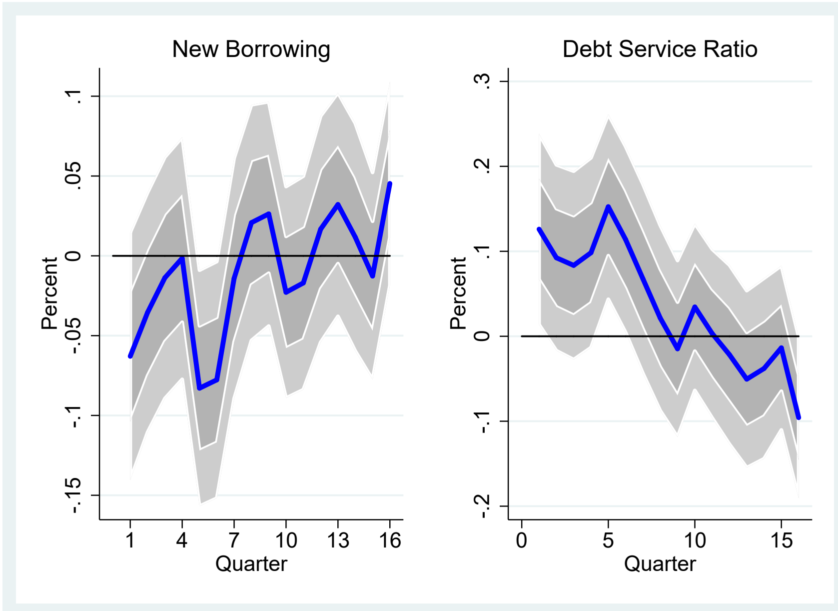

Figure 14: Impact of new borrowing and debt service on the probability of crises

As we can see from Figure 14, the results match our earlier findings: specifically, an increase in new borrowing decreases the chance of a crisis, while an increase in debt servicing increases the chance of a crisis. The extended crisis series, therefore, has no impact on the results.

We also check whether the definition of a crisis matters for our horse race of new borrowing and debt service ratio versus other more traditional credit measures. Table 4 suggests that the definition does not matter. The results match earlier findings in that new borrowing and debt service ratio play an outsized role compared with traditional credit measures when it comes to predicting crises. The coefficients for each are remarkably consistent across specifications.

Table 4: The effects of different credit measures on crisis probability (OLS)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
New Borr	-0.085***					-0.109***	-0.129***	-0.117**	-0.075**
	(0.032)					(0.041)	(0.046)	(0.049)	(0.033)
DSR	0.136**					0.164**	0.157**	0.176**	0.123**
	(0.055)					(0.067)	(0.061)	(0.076)	(0.051)
Credit Gr		0.006				0.048			
		(0.028)				(0.038)			
Cr/GDP Gr			0.002				0.025*		
			(0.008)				(0.013)		
Gap HP1				-0.010				0.025	
				(0.012)				(0.020)	
Gap HP2					-0.029				-0.019
					(0.017)				(0.017)
Observations	117	117	114	117	117	117	114	117	117
R^2	0.14	0.08	0.09	0.09	0.11	0.15	0.17	0.15	0.15

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
Estimated with baseline controls as in local projections figures (no lags).
DSR = debt service ratio, Credit Gr = quarterly credit growth, Cr/GDP Gr = year over year growth in credit-to-GDP ratio, Gap HP1 = one-sided Hodrick-Prescott filter deviation of credit-to-GDP from trend, Gap HP2 = two-sided Hodrick-Prescott filter deviation of credit-to-GDP from trend.
Source: Authors' calculations.

As a final robustness check, we focus on our vulnerabilities barometer, and by extension the importance of the debt-service ratio, by replacing new borrowing and the debt-service ratio in Tables 1 and 2 with the vulnerabilities barometer itself. We use the one-year-ahead dependent variable since the vulnerabilities barometer is unlikely to be of much use one quarter ahead. Table 5 indicates that the vulnerabilities barometer does a nice job of predicting a drop in economic growth, as expected. Similarly, as indicated in Table 6, it does a nice job of forecasting financial crises one year out.

Table 5: The effects of different credit measures on output growth

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Vulnerability Barometer	-0.332**					-0.382**	-0.465***	-0.459***	-0.532***
	(0.143)					(0.164)	(0.157)	(0.137)	(0.134)
Credit Gr		-0.036				0.188			
		(0.196)				(0.214)			
Cr/GDP Gr			-0.018				0.150**		
			(0.065)				(0.071)		
Gap HP1				0.095				0.244***	
				(0.086)				(0.090)	
Gap HP2					0.582***				0.682***
					(0.109)				(0.097)
Observations	114	114	111	114	114	114	111	114	114
R^2	0.34	0.30	0.26	0.30	0.52	0.35	0.33	0.37	0.63

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
Estimated with baseline controls as in local projections figures (no lags).
Credit Gr = quarterly credit growth, Cr/GDP Gr = year over year growth
in credit-to-GDP ratio, Gap HP1 = one-sided Hodrick-Prescott filter deviation of credit-to-GDP from trend,
Gap HP2 = two-sided Hodrick-Prescott filter deviation of credit-to-GDP from trend.
Source: Authors' calculations.

Table 6: The effects of different credit measures on crisis probability (OLS)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Vulnerability Barometer	0.055*					0.079**	0.073*	0.075**	0.070**
	(0.032)					(0.033)	(0.040)	(0.033)	(0.034)
Credit Gr		-0.043*				-0.089***			
		(0.023)				(0.032)			
Cr/GDP Gr			0.006				-0.020*		
			(0.009)				(0.011)		
Gap HP1				-0.014				-0.038***	
				(0.011)				(0.012)	
Gap HP2					-0.038**				-0.051***
					(0.016)				(0.018)
Observations	114	114	111	114	114	114	111	114	114
R^2	0.33	0.28	0.09	0.27	0.31	0.39	0.19	0.37	0.40

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
Estimated with baseline controls as in local projections figures (no lags).
Credit Gr = quarterly credit growth, Cr/GDP Gr = year over year growth
in credit-to-GDP ratio, Gap HP1 = one-sided Hodrick-Prescott filter deviation of credit-to-GDP from trend,
Gap HP2 = two-sided Hodrick-Prescott filter deviation of credit-to-GDP from trend.
Source: Authors' calculations.

8 Conclusions

The Canadian evidence supports the idea that debt service ratios have more predictive power than debt ratios for both aggregate spending and for financial crises. Our financial vulnerability barometer provides advanced warning of crisis periods and avoids false positives. It is therefore a useful tool for forecasting and policy analysis.

An immediate policy conclusion is that empirical models of financial system vulnerabilities (such as Pasricha et al. 2013, Duprey and Roberts 2017 and Duprey and Ueberfeldt 2018) should put more weight on the evolution of debt service ratios and other indicators related to debt servicing.

Further work is required on these indicators. More granular information on debt maturity and amortization rates would be useful. Households with different characteristics may have different propensities to borrow and consume. Households also differ in terms of the constraints they face on borrowing: when these constraints are severe, households may not be able to compensate for unexpected increases in their debt service burdens by additional borrowing.

There are also policy implications for the conduct of monetary policy. If monetary policy is to take financial stability concerns into account, central banks must be concerned with the dynamic effects of their policies: they will have to trade off the immediate effects on output and inflation with the future effects that work their way through the debt service channel.

Regulators concerned with the housing market and overall financial stability should also look beyond traditional credit-to-GDP measures and closely monitor the behaviour of debt servicing. This will allow a more accurate assessment of risk.

Our data sample covers a period when world interest rates have trended downward. This trend explains why debt ratios in Canada have increased over time: it is rational to borrow more to purchase productive assets when the costs of borrowing are low. If interest rates were to return to the levels last seen forty years ago, this would likely also mean a long-run reduction in debt ratios. Such decreases should not be interpreted as a sign of increasing financial stability since it might also be accompanied by temporary spikes in the household debt-service ratio if interest rates increase faster than levels of debt can be reduced. Monitoring the debt-service channel would be of crucial importance in such a context. However, yield curves in developed economies show no signs of anything more than modest increases in interest rates.

Data Appendix

The variables used for new borrowing and debt servicing are defined in the text.

The independent variables from Sections 4 and 5 and the data sources are defined as follows.

1. GDP growth: quarterly annualized growth rate of expenditure-based real GDP, seasonally adjusted. Source: Statistics Canada.

2. Real short rate: 3 month prime corporate paper rate minus inflation. Source: Statistics Canada and authors' calculations.

3. Lending spread: Prime business lending rate minus inflation. Source: Statistics Canada and authors' calculations.

4. Change in average lending rate on household debt: implied interest rate based on Statistics Canada household credit data and Statistics Canada interest payments (authors' calculations)

5. Real property price growth: quarterly growth rate (authors' calculations), where real property prices are calculated using nominal residential property prices from the BIS, and CPI comes from Statistics Canada.

The data underlying our vulnerability index are defined as follows:

1. Household debt to disposable income: monthly credit measured, turned quarterly (using month-end), divided by disposable income. Source: Statistics Canada, and authors' calculations.
2. Housing price to rent ratio: Indexed to 100. Source: OECD.
3. Non-financial corporate debt to GDP: quarterly, market value of debt securities, loans, accounts payable. Source: Statistics Canada.
4. Financial Institution debt to GDP: market value, Chartered and quasi-banks: debt securities, loans, accounts payable. Source: Statistics Canada.
5. Household debt to GDP: monthly credit measured, turned quarterly (using month-end), divided by GDP. Source: Statistics Canada, and authors' calculations.

Technical Appendix

The vulnerabilities barometer is based on what is known as "receiver operating characteristic" (ROC) analysis. In this appendix, we summarize some of the elements of the theory behind ROC analysis.[21]

Consider an indicator variable of financial distress that can take on a continuous range of values. If it exceeds a threshold value (we discuss the optimal choice of the threshold

[21] For further details, see Fawcett (2006), Flach (2010), Pasricha et al. (2013), Swets et al. (2000), and van Erkel and Pattynama (1998).

below), it signals potential financial distress within the next 24 months. Given historical data on the indicator variable, the observations can be divided into true positives, false positives, true negatives, and false negatives, as in the matrix below.

	Stress occurs	No stress occurs
Positive signal	True positive (A)	False positive (B)
Negative signal	False negative (C)	True negative (D)

An ROC curve measures the proportion of false positives $\left(\frac{B}{B+D}\right)$ on the horizontal axis and the proportion of true positives $\left(\frac{A}{A+C}\right)$ on the vertical axis for each possible value of the threshold. With a very high threshold, one is at the lower left-hand corner: there are no positives, so the true positive ratio is zero and the false positive ratio is zero. With a very low threshold, one is at the upper right-hand corner. At this corner all test results are positive, so that the true positive ratio is one, as is the false positive ratio. An ROC curve traces out the proportion of false positives (horizontal axis) against the proportion of true positives (vertical axis) as one varies the threshold for a positive "diagnosis." A better test (or indicator in our context) will increase the proportion of true positives as one lowers the threshold while at the same time the proportion of false positives increases only slowly. This gives a more concave ROC curve, and therefore a higher area under the curve (AUROC). Better indicators have higher AUROC measures. An ideal indicator would have no false positives ($B = 0$) and no false negatives ($C = 0$), so it would be situated at the upper left-hand corner of the ROC graph.

A shortcut measure of the quality of an indicator is its "signal-to-noise" ratio, given by

$$\text{SNR} \equiv \frac{A/\left(A+C\right)}{B/\left(B+D\right)}.$$

A value greater than one means that an indicator performs better than a coin flip.

This still leaves the problem of picking the optimal value of the threshold. Clearly this would depend on the relative costs of false positives versus false negatives. According to van Erkel and Pattynama (page 93):

> Intuitively, one would identify the 'optimal' operating point as the point on the ROC curve that is closest to the ideal upper left-hand corner. Determining the optimal operating point on the ROC curve, however, involves both clinical and financial issues. For instance, pneumonia is a disease in which a large therapeutic gain can be achieved at relatively little cost and with few complications of the antibiotic treatment. Thus, a certain amount of false positive results is acceptable, whereas false negative results are less desirable. As a result, in testing for pneumonia, the threshold value is usually set low. In other words, the optimal operating point will move towards the upper right-hand part of the ROC curve. On the other hand, when we are considering a costly and potentially harmful treatment with only little therapeutic benefit, it is the false positive results that are to be limited. The optimal range of the operating point will thus, shift towards the lower left-hand corner of the ROC graph. Clearly, in determining the optimal threshold value, we have to take into account all the clinical and financial consequences of the different test results. Ideally, such decisions should be made by linking the constructed ROC curve to explicit decision analysis.

In the absence of a way of parameterizing the costs of false positives versus false negatives, we follow Pasricha et al. (2013) and minimize the following objective function in order to generate the threshold. This objective function can be thought of as an "impurity measure", or, the extent to which signals mis-classify stress and tranquil periods:

$$\min_{\tau_i} \left(\frac{1}{A+B+C+D} \left(\frac{CD}{C+D} + \frac{AB}{A+B} \right) \right).$$

With no false negatives or false positives ($B = C = 0$), the value of this function would be zero.

References

Ambler, Steve and Jeremy Kronick and Farah Omran. 2019. "Household Insolvency." *Intelligence Memo* C.D. Howe Institute, April 30

https://www.cdhowe.org/intelligence-memos/ambler-kronick-omran-household-insolvency

Congdon, Tim. 2017. "False Alarm." *EA Magazine* issue 9, 44–46.

https://iea.org.uk/wp-content/uploads/2017/03/EA-SPRING-2017_CITY-VIEW.pdf

Drehmann, Mathias and Mikael Juselius. 2012. "Do Debt Service Costs Affect Macroeconomic and Financial Stability?." BIS Quarterly Review. Bank for International Settlements. September.

Drehmann, Mathias and Mikael Juselius. 2014. "Evaluating Early Warning Indicators of Banking Crises: Satisfying Policy Requirements." *International Journal of Forecasting* 30, 759–780.

Drehmann, Mathias, Mikael Juselius and Anton Korinek. 2017. "Accounting for Debt Service: The Painful Legacy of Credit Booms." Working Paper 645, Bank for International Settlements.

www.https://www.bis.org/publ/work645.pdf

Duprey, Thibaut and Tom Roberts. 2017. "A Barometer of Canadian Financial System Vulnerabilities." Staff Analytical Note 2017-24, Ottawa, Bank of Canada.

https://www.bankofcanada.ca/2017/12/staff-analytical-note-2017-24/

Duprey, Thibaut and Alexander Ueberfeldt. 2018. "How to Manage Macroeconomic and Financial Stability Risks: A New Framework." Staff Analytical Note 2018-11, Ottawa, Bank of Canada.

https://www.bankofcanada.ca/2018/04/staff-analytical-note-2018-11/

Fawcett, Tom. 2006. "An Introduction to ROC Analysis." *Pattern Recognition Letters* 27, 861–874

Flach (2010), "ROC Analysis." in Claude Sammut and Geoffrey Webb (eds.), *Encyclopedia of Machine Learning*. Boston, MA, Springer, 869–875.

Jorda, Oscar. 2005. "Estimation and inference of impulse responses by local projections." *American Economic Review* 95, 161–182.

Pasricha, Gurnain, Tom Roberts, Ian Christensen and Brad Howell. 2013. "Assessing Financial System Vulnerabilities: An Early Warning Approach." *Bank of Canada Review* Autumn, 10–19.

Swets, John, Robyn Dawes and John Monahan. (2000) "Better Decisions through Science." *Scientific American* October, 82–87.

van Erkel, Arian and Peter Pattynama (1998), "Receiver Operating Characteristic (ROC) Analysis: Basic Principles and Applications in Radiology." *European Journal of Radiology* 27, 88–94.