The Bank of Canada's Analytic Framework for Assessing the Vulnerability of the Household Sector

Ramdane Djoudad

Introduction

Changes in household debt-service costs as a share of income—i.e., the debt-service ratio, or DSR—are a measure of changing risk associated with household debt. While aggregate data provide an indication of average shifts in household debt positions, such variations frequently obscure vulnerabilities that only a review of the microdata can reveal. The availability of microdata for this type of review has assisted the Bank in developing an analytical framework for assessing risk in the household sector.[1]

Although the DSR is not the only barometer of the financial health of households, it remains a good indicator of their vulnerability. A rise in the DSR, for example, increases the vulnerability of households to negative shocks and can also have potential adverse consequences for the balance sheets of financial institutions. Since household debt accounts for a significant proportion of the loan portfolios of banks, shifts in household vulnerability arising from potential variations in macroeconomic conditions must be monitored. This report outlines the Bank's framework for analyzing changes in household vulnerability as described in the June and December 2009 issues of the *Financial System Review* (FSR),[2] as well as recent improvements to that framework. The unique feature of the framework is the use of microdata in stress-testing simulations to measure the impact of various shocks (debt, interest rates, employment, amortization period, etc.) on the distribution of the DSR and, ultimately, on household solvency. These analyses are an attempt to gauge the impact of an adverse shock under simulated conditions rather than to identify the most likely changes in the financial conditions of households.

There are three steps in the stress-testing exercise **(Table 1)**. In Step 1, the key assumptions of a scenario representing a macro environment under stress are defined. The scenario should be consistent with the Bank's assessment of possible risks to the household sector. For example, in the December 2009 issue of the FSR, one of the main developments we wanted to evaluate was a continuation of strong credit growth in an environment of rising interest rates. Once the aggregate scenario is set (Step 1), we need to distribute the effect across individual households (Step 2). Finally, based on the evolution of the DSR distribution, we estimate the effects of an adverse shock on the credit losses at banks (Step 3).

Table 1: Steps in the stress-testing exercise

Step 1	Step 2	Step 3
• Establish the key assumptions for the macro scenario: – Growth in aggregate credit and income – Interest rate path	• Calculate the implications of the macro scenario for the distribution of the household debt-service ratio	• Estimate the impact of adverse shocks on bank loan portfolios

Two major improvements have recently been made to the methodology. First, those buying a home for the first time have been explicitly taken into account as a separate class in Step 2. Second, the risk assessments in Step 3 will be strengthened by combining elements from previous exercises reported in the June and December 2009 issues of the FSR. Specifically, household vulnerabilities will evolve over time by simulating changes in indebtedness and interest rates (as in the December 2009 FSR), and potential losses at banks will then be assessed using an explicit employment shock comparable to the one described in the June 2009 FSR.

1 Data are from the Canadian Financial Monitor (CFM) annual survey of approximately 12,000 households conducted by Ipsos Reid. The survey was launched in 1999.

2 *Financial System Review*, June 2009, pp. 21–23 and December 2009, pp. 23–26.

The Data

The DSR derived from microdata includes principal repayments on all instalment loans. To calculate the DSR, its three major components are evaluated: household debt, interest rates, and household income, as shown in the following formula.

$$\mathrm{DSR} = \frac{\sum \mathrm{Payments}}{\mathrm{Gross\ income}} = \frac{\sum(\mathrm{Principal} + \mathrm{Interest})}{\mathrm{Gross\ income}}. \quad (1)$$

The microdata used for the calculation include credit card debt, personal loans, personal lines of credit, vehicle loans, and mortgage loans. The following information is available for all loans except credit card debt:

- the amount of the monthly payment
- the effective interest rate
- the term[3] of a mortgage loan (in years), but not its maturity date
- the balance of the loan

In previous issues of the FSR,[4] the Bank reviewed the distribution of the DSR using microdata to better determine how debt was spread across various households classified by income. Determining the distribution of risk among households naturally requires a review of the upper tail of the DSR distribution since, all things being equal, households with a high DSR will obviously have a more difficult time meeting their financial obligations. Thus, the greater a household's debt load, the greater its sensitivity to idiosyncratic shocks, such as divorce or a serious illness, or to economic shocks, such as the loss of a job. A household's assets are also a significant factor in assessing its ability to weather negative financial shocks.

Chart 1 shows the cumulative distribution of the DSR for 2009.[5] This distribution indicates that the majority of households are below the critical 40 per cent threshold. Households with a DSR above this threshold may have difficulty meeting their debt obligations. By comparing this distribution with that of previous years, we can determine the changing profile of household sensitivity to shocks. However, a methodological framework is required to gauge the effect of certain shocks on the distribution. The purpose of this report is to describe the Bank's simulation model for quantifying the effects of changes in certain macroeconomic variables on the distribution of the DSR and, ultimately, on potential losses at banks.

3 Data are available for 6-month, 1-year, 2-year, 3-year, 5-year, 7-year, 10-year, and variable-rate mortgage loans.

4 See the five issues of the FSR published from December 2007 to December 2009.

5 For more details on the historical profile of changes to the DSR and the proportion of vulnerable households, see the December 2006 issue of the FSR, pp. 15–16.

Chart 1: In 2009, the majority of households had a debt-service ratio below 40 per cent
Cumulative distribution of the debt-service ratio, 2009

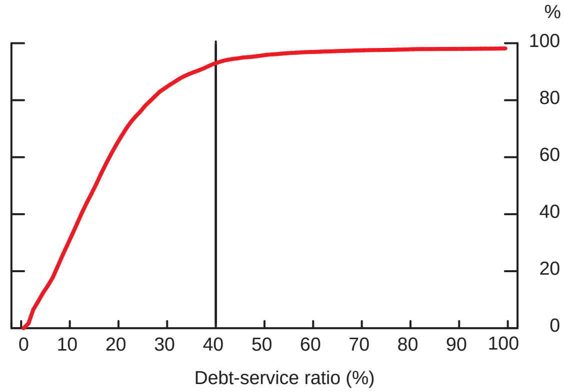

Source: Ipsos Reid

Impact on the Distribution of the Debt-Service Ratio

Having defined the aggregate hypothesis consistent with the macroeconomic scenario established in Step 1, the impact of these assumptions on the distribution of the DSR is assessed in Step 2. Simulations are carried out over a three-year horizon.

In this model, interest rate shocks affect only interest payable and not the amount of principal repayments. Consequently, interest payments must be distinguished from repayments of principal. The variable PC represents a household's total annual loan payments, SC is its current credit balance, and i, the applicable interest rate.

The following formula is used to determine the approximate amount of the principal repayments:

$$\mathrm{Principal} = PC - \mathrm{Interest} = PC - (SC^{*}\ i). \quad (2)$$

When simulations are performed, principal payments[6] are deemed to be a constant share of the credit balance:

$$\mathrm{Share\_Principal} = (\mathrm{Principal}/SC). \quad (3)$$

Thus, a household is required to make the following payment in each period:

$$PC = SC^{*}(\mathrm{Share\_Principal} + i). \quad (4)$$

6 In fact, the share of principal repayments may vary over time. However, since the simulations are performed over a short period of time, we do not think that this will significantly affect the results.

Future payments and the dynamics of the DSR will be determined by the simulated profile of changes in household income and debt, as well as in interest rates.

Interest rates

To design an interest rate scenario, we must define a profile of changes to the overnight rate (Step 1). For example, in the December 2009 issue of the FSR (pp. 23–24), the Bank considered two hypothetical paths for the overnight rate. The first was a reflection of market expectations embodied in current yields on Government of Canada securities, while the second assumed a sharper rise. Additional assumptions are required for the profiles of risk and term premiums on household debt. In the December 2009 FSR (p. 24), the Bank assumed that risk premiums would return to their historical levels at the end of the simulation period.[7]

Since we know the date on which each household completed the Ipsos Reid questionnaire, we are able to calculate the risk premium on variable-rate loans by subtracting the overnight rate from the actual interest rate. It is assumed that households make credit card payments equal to 2 per cent of the monthly balance; i.e., the minimum payments generally required by card issuers. It is also assumed that variable interest rates apply to all other types of consumer loans (personal loans, personal lines of credit, and vehicle loans).[8] Variable-rate debt responds immediately to changes in the overnight rate.

For simplification, we assume that the proportion of households whose mortgages are renewed in a given year is equal to the reciprocal of the term to maturity. For example, for a 5-year term, 20 per cent (1/5 = 0.2) of households would renew their mortgage each year (5 per cent per quarter).

Heterogeneity in income growth

Income is the second variable required to plot the projected evolution of the DSR. The approach used was to divide households into five classes, based on income (for details, see Djoudad 2009). The following equation represents the distribution of income growth for a particular class:

$$\text{Income} \sim N\left(r_j,\ \sigma_j\right) \quad j = 1, 2, 3, 4, 5, \qquad (5)$$

where

j = household income class

r_j = average income growth of households in class j

σ_j = estimated standard deviation of income growth for households in class j (Djoudad 2009).

In this framework, income growth is assumed to be heterogeneous within each class. Between classes, the mean and standard deviation may be assumed to be similar or different, although overall growth must be consistent with the aggregate scenario (Step 1). For example, a shock to income (Step 1) may have a greater impact on the income growth of households in the lowest income classes (1 and 2) than for households in the highest income classes (3, 4, and 5).

Heterogeneity in the growth of household debt

The macroeconomic scenario considered includes assumptions for total growth in mortgage and consumer debt. That said, all households cannot be presumed to experience identical debt growth. The distribution of the growth of aggregate debt across income classes must therefore be determined. Since all households are not comparable, the simulation model incorporates household heterogeneity by allowing the growth of each household's debt to depend on its specific socioeconomic characteristics and certain empirical relationships (as described below). A specific distinction is made between first-time homebuyers, who have yet to contract mortgage debt, and all others.

First-time homebuyers

First-time homebuyers have accounted for a significant share of the growth in mortgage credit in recent years. According to some analysts (e.g., the Canadian Association of Accredited Mortgage Professionals (CAAMP 2010)), nearly 50 per cent of all homebuyers were new to the market in 2009. After purchasing their first home, their debt exceeds the average for Canadian households. While first-time homebuyers were implicit in previous exercises, they are now taken explicitly into account in the model. They must therefore be distinguished from other households to avoid unduly increasing the debt loads of current mortgage holders, thus inflating the proportion of vulnerable households. Taking first-time homebuyers into account leads to lower levels in the measures of vulnerability, given that a significant share of new mortgages goes to households that previously had no mortgage debt.

To illustrate the impact of the new methodology, **Table 2** compares simulation results of a model with and without first-time homebuyers. These simulations use updated data for 2009H2 and 2010Q1. Under Scenario 2 of the December FSR, the results indicate that taking explicit account of first-time homebuyers lowers the proportion of vulnerable households to 7.4 per cent from 8.4 per cent, and the percentage of debt owed by these vulnerable households to 14.3 per cent from 17.2 per cent, by the end of 2012Q2.[9]

In each period, new households that have neither taken on a mortgage nor purchased a home are drawn from our data on

7 The methodology is flexible and lends itself to a variety of scenarios.

8 Credit cards are at fixed rates; personal lines of credit account for almost 75 per cent of all remaining consumer loans, most of which are at variable rates.

9 The numbers reported in Table 2 for the previous methodology differ from those reported in the December 2009 FSR, owing primarily to a correction to the program code.

Table 2: Impact on the vulnerability measures of introducing first-time homebuyers (%)

Period	Previous methodology		Explicitly taking into account first-time homebuyers	
	Proportion of households with DSR > 40%	Proportion of debt owed by households with DSR > 40%	Proportion of households with DSR > 40%	Proportion of debt owed by households with DSR > 40%
2010Q1	5.1	9.7	5.0	9.6
2010Q4	5.6	11.0	5.4	10.5
2011Q4	7.6	15.2	6.8	12.9
2012Q2	8.4	17.2	7.4	14.3

households and are added to the sample of homeowners. According to data from the CAAMP, in 2007 the average gross DSR for all new mortgage borrowers was around 23 per cent. Households that are added to the sample are assigned a share of new debt on the basis of their income and the distribution of the DSR for first-time buyers, consistent with the observed distribution in recent years.[10]

Other households

The difference between the aggregate new debt and the portion that has been attributed to first-time homebuyers is the residual debt, which is the share of mortgage debt attributed to other households. First, the share of mortgage debt incurred by first-time homebuyers is subtracted from aggregate debt growth (total debt and mortgage debt). Next, residual debt growth is spread among those households with previously incurred debt. For example, if the scenario assumes a 10 per cent increase in aggregate mortgage debt, half of which was taken on by first-time homebuyers, the mortgage debt of all other households whose homes are already mortgaged should increase by only 5 per cent.

Having determined interest rates, increases in real estate prices, and the rate of income growth for individual households, as well as their DSRs, we calculate average growth rates of total credit and mortgage credit, using equations (6) and (7). With the exception of the DSR, all other variables are expressed in first differences.

$$\text{Total household credit} = \text{F(income, interest rates, housing wealth, dsr)}. \tag{6}$$

$$\text{Mortgage credit} = \text{F(income, interest rates, housing wealth, dsr)}. \tag{7}$$

The empirical relationships described in equations (6) and (7) serve to determine changes in the total-credit and mortgage-credit profiles of each of the remaining households. In these equations, the growth of debt depends on household characteristics and the assumptions underlying the macroeconomic scenario. To spread this residual debt among households with previously incurred debt, we use equation (5) to generate a stochastic distribution of income for all households. Equations (6) and (7) were estimated on the basis of data[11] pertaining to various classes of households, while taking into account such variables as the household's labour market status, level of education, place of residence, family income, and housing wealth, as well as interest rates. A detailed analysis of estimation results is provided by Djoudad (2009) and Dey, Djoudad, and Terajima (2008).

Most financial institutions consider a DSR of 40 per cent to be the threshold above which a household may have difficulty making loan payments. Hence, it is more difficult for households with a DSR of 40 per cent or more to incur additional debt, since financial institutions will scrutinize their loan applications more closely. Such households therefore find themselves with greater constraints, and so we surmise that their debt behaviour changes as they reach the threshold. As a result, the model allows the marginal effect of a rise in income or interest rates on debt to diminish as the household reaches a DSR threshold of 40 per cent.

Determining the evolution of the DSR

Now that we have established how interest rates, income, and debt evolve, we are able to recalculate each household's debt-service costs based on the interest rate applicable to any new or renewed debt. We use an individual household's payment schedule and income to calculate changes to its DSR profile over the entire simulation period. These household-specific results are used to calculate the distribution of the DSR across all households.

Risk Assessment

Past issues of the FSR have reported two types of DSR simulation exercises that assess the impact of changes in macroeconomic conditions on the financial health of households and, ultimately, the balance sheets of financial institutions.

Under the first type of simulation,[12] the Bank assessed the medium-term risks stemming from increasing indebtedness in an environment of rising interest rates. While this type of exercise does result in a vulnerability metric (i.e., the share of households where the DSR is equal to or greater than a critical threshold), it provides no direct measure of the losses financial institutions are liable to sustain.

In the June 2009 issue of the FSR, the Bank attempted to assess the impact of a more severe negative shock on the

10 The information on mortgage terms for first-time homebuyers is taken into account in calculating the monthly maturities. Although the methodology can accommodate alternative scenarios, all first-time homebuyers are assumed to have a 5-year mortgage.

11 Specifically, CFM survey data for the years 1999 to 2007 were used.

12 See the December 2009 issue of the FSR, pp. 23–26, for an example.

Canadian economy than was anticipated at the time, by introducing an explicit macroeconomic shock to employment. This exercise (unlike the type of simulation described above) provided a direct assessment of the impact of potential losses on the balance sheets of financial institutions. However, debt, income of the employed, and interest rates were assumed to be constant. These were reasonable simplifying assumptions, since the purpose of that exercise was to assess near-term risks, but they would not be realistic for assessing risks over a longer horizon.

The impact of any negative shock on the balance sheets of households and, ultimately, on those of financial institutions depends on the significance of the vulnerabilities at the time the shock occurs. Accordingly, future stress tests will combine the basic features of both types of simulation within our framework. The effect of changes in income, debt, and interest rates on the DSR distribution will be simulated, and the distributions generated for each time horizon will be used to evaluate the impact of hypothetical shocks to employment on the loan losses of financial institutions. This approach should support a more sophisticated analysis of how risk is transferred from households to the financial system.

Employment shock

A negative shock to employment would result in a significant loss of income for households that are affected. In our model, the distribution of job losses among sampled households is random (retirees, students, and other households with no employment income would not be affected).[13] Sources of funds for unemployed households would be limited to employment insurance, provided they are eligible, and any liquid assets they may have (balances in chequing and savings accounts, term deposits, GICs, etc.). It is possible that illiquid assets could be sold and included in the funds available to households. However, in a systemic crisis, households may have difficulty selling off their assets without triggering a significant drop in prices. The price declines would exacerbate the financial stress. If a broader range of assets were used, then the second-round effects would also need to be considered in the model. Overall, restricting the calculation to liquid assets should not bias the conclusions.

According to empirical data, only a fraction of households would be eligible for employment insurance benefits in the event of a job loss. Given that all households have fixed expenses (housing, food, etc.), it is assumed that half of the funds available to a household would be used for such expenses and would not be available to cover debt-service costs. We determine a household's ability to fulfill its financial obligations by comparing available funds (including liquid assets) to total payment requirements over the period of unemployment. The longer the period of unemployment lasts, the lower will be the remaining resources available to the household to meet its debt-service obligations and the higher the probability of it becoming insolvent. If a household is unable to meet its debt obligations for more than three consecutive months, it is considered insolvent and its unsecured outstanding debt is considered a loss to financial institutions.

The average period of unemployment is a critical factor in assessing whether a household will become insolvent. Consistent with historical evidence, the higher the unemployment rate, the longer the average period of unemployment will be. Our simulations assume that the duration of unemployment varies among households, following a chi-squared distribution.

The impact of a shock on the default rate

Our measures of vulnerability include the share of household income required to cover debt-service costs. In our estimation, households that devote more than 40 per cent of their income to servicing debt are far more vulnerable to shocks than those carrying a lighter debt load. The proportion of vulnerable households and their share of debt are measures of household vulnerability to external (economic or personal) shocks. These vulnerability measures are a useful summary statistic that is often reported in our stress tests, but they do not represent a direct measure of losses when a shock is realized.

To assess the impact of a shock on the financial system, we estimate the likely number of households that would be unable to meet their payment obligations in the event of a shock. In the June 2009 FSR, the Bank used the method set out above to determine the proportion of households that would become insolvent given a rise in the unemployment rate, as well as the share of debt incurred by such households. Based on these results, the share of unsecured debt owed by these households is calculated to estimate the losses that banks are liable to incur and their impact on Tier 1 capital (equation 8). Unsecured debt does not include mortgage loans, secured lines of credit, and other secured consumer loans.[14]

$$\text{Adjusted Tier 1 capital ratio} = \frac{\text{Tier 1 capital} - \text{Losses on unsecured loans}}{\text{Risk-weighted assets} - \text{Losses on unsecured loans}}. \qquad (8)$$[15]

Conclusion

Microdata are a valuable source of information for assessing the risks associated with household debt. The Bank of Canada has been using microdata for several years as a

13 A future research objective is to adjust this distribution to stylized facts. We may assume, for example, that a negative shock to employment will have the greatest impact on low-income or younger workers.

14 Mortgages are excluded, since about half are insured, while the rest have a low loan-to-value ratio.

15 Levels of capital are assumed to increase at some rate before the shock occurs.

complement to its analysis based on aggregate data. This report presents methodological advances made by the Bank in using these data.

Examples of the shocks considered here demonstrate the possible applications of this framework. Of course, this type of tool continues to evolve and could be enhanced by a more refined representation of the economic behaviour of households. For example, certain random data draws could be governed by behavioural rules more in line with economic theory and the stylized facts. Our estimations of the parameters, by household class, using equations (6) and (7), are a step in that direction. We are currently enhancing the model by fleshing out the links between household characteristics and measures of vulnerability. There is also a need to refine the way income is determined. For income growth (equation 5), for example, we could estimate a structural equation.

Although this model is a simplified version of the real world, it nonetheless provides an innovative and promising means of studying household vulnerabilities and risks to the banking system. It is a flexible empirical tool that can be adapted to take into account a wide variety of alternative scenarios.

References

Canadian Association of Accredited Mortgage Professionals (CAAMP). 2010. *Revisiting the Canadian Mortgage Market—Risk Is Small and Contained*, January 2010.

Dey, S., R. Djoudad, and Y. Terajima. 2008. “A Tool for Assessing Financial Vulnerabilities in the Household Sector.” *Bank of Canada Review* (Summer): 45–54.

Djoudad, R. 2009. “Simulations du ratio du service de la dette des consommateurs en utilisant des données micro.” Bank of Canada Working Paper No. 2009–18.