An Improved Framework for Assessing the Risks Arising from Elevated Household Debt

Umar Faruqui, Xuezhi Liu and Tom Roberts

Introduction

Since 2008, the Bank of Canada has used a micro-simulation model as one of its tools to assess the risks to financial stability emanating from the elevated debt burdens of Canadian households. The strength of this approach is its use of actual household balance sheets to examine the distribution of debt within the household sector.[1] Unlike aggregate measures such as the ratio of household debt to income, this distributional information provides insight into the most vulnerable segments within the household sector—where problems would first arise.

This report describes the improvements that have been made to address limitations in the previous version of the model and incorporates a set of scenarios to illustrate these improvements. Most notably, the model can now simulate a multi-year unemployment shock, with expanded dynamics that include the asset side of household balance sheets, in addition to the debt and income dynamics that already made up the core of the model (**Box 1**).[2]

These features provide increased flexibility to adjust the characteristics of a shock scenario. The enhancements allow for greater insight into the interactions between different sources of risk, as well as a better assessment of the evolution of risk over the simulation horizon, as indicated by the share of households with an elevated debt-service ratio (the ratio of debt payments to income) and the effect of a shock on loan arrears.

The new version of the model—the Household Risk Assessment Model (HRAM)—will be used to conduct the household stress tests reported in this and future issues of the *Financial System Review* (FSR). Although a single model cannot provide a comprehensive account of all possible risk interactions, HRAM is an important part of the Bank's ongoing development of complementary approaches for monitoring risks in the household sector.

More Flexible Modelling of Unemployment Shocks

A negative shock to the labour market is an important potential source of financial stress in the household sector. Previous analyses published in the FSR examined this risk by first simulating the distribution of the debt-service ratio (DSR) under a stable set of macro assumptions for 10 quarters, and then imposing an unemployment shock on this distribution. The model then calculated loan arrears over a 1-year period following the unemployment shock. HRAM, however, offers greater flexibility to assess the impact of a wide range of scenarios on household vulnerabilities and loan arrears by adjusting the timing, severity and persistence of the shock. For example, it can evaluate the impact on loan arrears of a sharp but short-lived unemployment shock or a slow and persistent increase in unemployment over a multi-year horizon. By allowing all of the scenario assumptions to evolve simultaneously with the unemployment shock, HRAM also increases the coherence of the scenario throughout the stress test.[3]

The number of unemployed households in each period is determined by the assumptions for the unemployment rate and its average duration (in weeks). Unemployed households in a given period are defined as the sum

1 Data are from the Canadian Financial Monitor survey compiled by Ipsos Reid. The data include information on balance sheets, income, debt payments and other characteristics related to household finances for about 12,000 households that are representative of the Canadian population.

2 The version of the model used in previous analyses reported in the *Financial System Review* is outlined in Djoudad (2010; 2012).

3 For example, the shock scenario described later in this article incorporates adjustments to credit growth, income growth and risk premiums in response to a macroeconomic downturn. In the previous version of the model, there was no adjustment of these other variables after the unemployment shock was introduced at the end of the simulation period.

Box 1

Core Features of the Simulation Model

There are three steps in the stress-testing exercise. First, a scenario representing a stressed macro environment is defined. This scenario includes assumptions for the level of interest rates and the aggregate growth rates of household credit and income. Second, using the latest microdata as the starting point, each household's debt-service ratio (DSR) is simulated over the projection horizon by allocating credit and income growth across individual households.[1] While these individual growth rates vary, the combined outcomes are consistent with the assumptions in the aggregate scenario. Finally, based on the simulated distribution of the DSR across all households, we estimate the effects of an adverse shock on loans in arrears.

The core debt and income dynamics in HRAM are unchanged from the previous version of the model. The distribution of credit growth depends on household-specific characteristics such as income, housing wealth, a household's initial DSR and whether it is a first-time homebuyer. Consistent with empirical evidence, households with higher current income and wealth and a lower DSR tend to have better access to credit and can accumulate more debt than other households. Unemployed households cannot obtain additional credit.

Income growth is simulated by grouping employed households into five income categories. For each category, household-specific income growth is randomly drawn from a normal distribution. Mean income growth and the standard deviation can vary across income groups, in line with the empirical evidence (Djoudad 2012).

The borrowing cost of variable-rate debt responds immediately to changes in the overnight rate or the risk premium for household debt. For fixed-rate mortgages, 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 mortgages each year at current rates. Thus, the average borrowing rate for all outstanding debt—the "effective household borrowing rate"—will change more gradually than the interest rate for new loans.

1 In the new version of the model, households can also accumulate financial assets.

of households that were unemployed in the previous period and remain unemployed, plus newly unemployed households.[4] Once a household is unemployed, it loses its labour income but may receive employment insurance. To account for living expenses, only a fixed proportion of employment insurance, and of any income from a second income earner, are available to unemployed households for debt-servicing costs. If these sources of funds are insufficient to cover the household's needs, it can draw on its liquid financial assets[5] until all of these assets are depleted.[6] When a household becomes re-employed, it is reassigned its initial employment income.[7]

For each period, HRAM keeps track of every household's labour force status, ability to make debt payments and financial wealth. Once a household is three months or more behind in its debt payments, it is categorized as being in arrears.[8] A distinguishing feature of the unemployment shock in HRAM is that, unlike the previous version of the model, households can cycle in and out of periods of unemployment over the simulation.

Improved Dynamics for Household Financial Assets

The financial assets held by households play a significant role in the model because unemployed households can use these assets to service their debt and avoid (or delay) going into arrears. As a result, a downturn in financial markets or changes in savings behaviour could affect arrears. To capture these effects, the new model allows the asset side of household balance sheets to evolve over the simulation period.

Changes in the total value of financial assets occur through three channels in HRAM: (i) household savings from current income; (ii) asset-price movements; and (iii)

4 Unemployment is assigned randomly across households in the workforce. In reality, some households (based on region, age and employment sector) would have a higher probability of unemployment than others, which might become important if vulnerable households are disproportionately affected by a downturn. The duration of unemployment for each unemployed household is also assigned at random and is limited to a range of one to 99 weeks. The assumed distribution of duration reflects empirical evidence that, while periods of unemployment are generally short, there is a significant share of long-term periods. For example, among the unemployed households questioned by Statistics Canada in 2011 for the Labour Force Survey, about 62 per cent had been unemployed for 13 weeks or less, while 22 per cent had been unemployed for more than 26 weeks.

5 These include cash, bonds, stocks and mutual funds.

6 In the model, households cannot borrow to cover a shortfall in debt-servicing obligations.

7 In reality, human capital tends to depreciate over longer periods of unemployment, which would affect a household's income when it becomes re-employed. Tests suggest that this modification would likely not have a significant impact on the simulation results.

8 Financial institutions typically start to provision for losses when loans are three months in arrears.

household dis-savings (for example, when liquid assets are drawn down during periods of unemployment). The first two channels were not present in the previous version of the model, and the third was confined to the end of the simulation exercise when the unemployment rate was assumed to rise.

The aggregate savings rate of the household sector follows an assumed path for the scenario, but the savings rate varies across households. Specifically, households are allocated savings based on an econometric estimation of the relationship between savings and household-specific characteristics such as income, debt payments, the current level of assets (divided by the number of working years) and age. The savings allocation for any particular household cannot exceed a budget constraint (which takes into account debt payments). This approach helps to ensure that savings are allocated according to financial capacity and savings preferences, and reflects empirical evidence that savings behaviour is uneven across households.

Changes in the prices of financial assets have a direct effect on the wealth of households. The returns on key categories of assets (e.g., stocks and bonds) are chosen to be consistent with the scenario. Household asset holdings are adjusted at the end of each simulation period to reflect both savings and the returns on financial assets. During a recession, when unemployment rises and returns on financial portfolios are negative, a household's wealth—and potentially its ability to weather a period of unemployment—declines. Simulation results presented later in this report suggest, however, that the asset-price channel has a smaller *direct* impact on loan arrears than the savings channel.

As noted earlier, liquid financial wealth in the form of cash and marketable securities is used by unemployed households to make debt payments if employment insurance and other income (for two-income households) are insufficient. Thus, for some unemployed households, financial wealth declines from one period to the next because of withdrawals for debt payments.

Putting HRAM to Work: Outputs, Assumptions and Analysis of the Key Changes to the Model

In this section, we outline the impact of the changes described above on the results produced by the model. The analysis is presented in two steps. First, we review the main outputs of the model and describe the scenario assumptions for our illustrative exercise. Second, we assess the results.

Model outputs and scenario assumptions

Like the previous version of the model, HRAM has three main simulation outputs: (i) the share of vulnerable households among indebted households, where a vulnerable household is defined as having a DSR that is equal to or greater than 40 per cent;[9] (ii) the share of total household-sector debt held by *vulnerable* households; and (iii) the share of total household debt that is three or more months in arrears.[10, 11]

The first two measures show how the most vulnerable part of the distribution of household DSRs changes over the simulation period. The loan arrears rate describes the share of outstanding household debt that is currently distressed.

To illustrate HRAM's flexibility, we consider the model's response under four scenarios: a control scenario, a short-lived unemployment shock, and a persistent unemployment shock *with* and *without* asset dynamics from savings and returns on financial assets.[12] A subset of these assumptions is highlighted in **Chart 1**, **Table 1** and **Chart 2**.

Chart 1: Unemployment rate: Assumptions

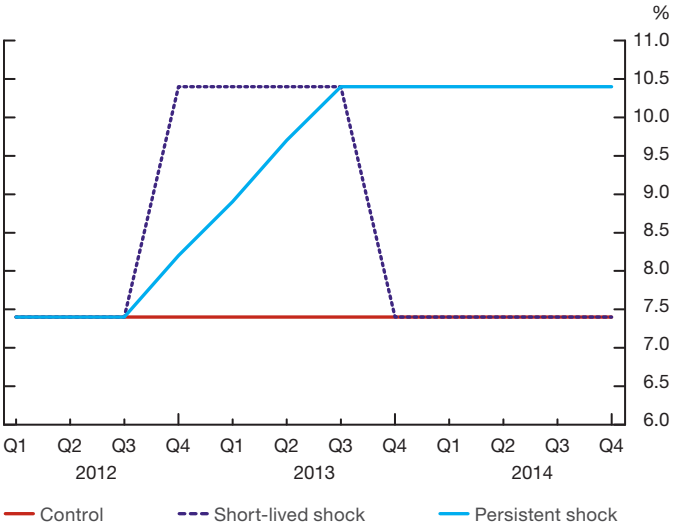

9 Following industry standards, a household is expected to have more difficulty making loan payments when its DSR is equal to or greater than 40 per cent.

10 These measures are described in more detail in Djoudad (2010).

11 The model generates a flow of new arrears in each simulation quarter, while available data on arrears are measured as a stock (as a percentage of debt). To translate the flow into a stock, the model assumes that new mortgage arrears are written off by banks (or return to regular payments) after 2.5 quarters, on average, whereas, consistent with financial reporting rules, the duration is one quarter for consumer loan arrears (which exclude mortgages).

12 None of these scenarios represents the Bank's view on the most probable outcome for the macroeconomic environment or household sector risk. Instead, they illustrate vulnerabilities in the household sector under a range of hypothetical situations.

Table 1: Main assumptions for control and shock scenarios

	Credit and income (quarter-over-quarter annualized growth, per cent)				Interest rates (basis points)				Unemployment (per cent)	
	Control		Persistent shock		Control		Persistent shock		Control	Persistent shock
	Growth of total household credit	Growth of disposable income	Growth of total household credit	Growth of disposable income	Overnight rate[a]	Effective household borrowing rate[b]	Overnight rate	Effective household borrowing rate[b]	Rate	Rate
2012Q1	4.7	3.9	4.7	1.1	100	490	100	493	7.4	7.4
2012Q2	5.3	3.8	5.3	3.8	100	473	100	477	7.4	7.4
2012Q3	5.3	3.8	5.3	3.8	109	461	100	462	7.4	7.4
2012Q4	1.0	3.8	1.0	0.0	122	465	100	497	7.4	8.2
2013Q1	1.0	3.8	1.0	0.0	131	467	100	533	7.4	8.9
2013Q2	1.0	3.8	1.0	0.0	143	471	100	569	7.4	9.7
2013Q3	1.0	3.8	1.0	0.0	148	472	100	607	7.4	10.4
2013Q4	2.5	3.8	2.5	2.0	152	473	100	609	7.4	10.4
2014Q1	2.5	3.8	2.5	2.0	155	473	100	611	7.4	10.4
2014Q2	2.5	3.8	2.5	2.0	158	473	100	613	7.4	10.4
2014Q3	2.5	3.8	2.5	2.0	160	473	100	614	7.4	10.4
2014Q4	2.5	3.8	2.5	2.0	162	472	100	616	7.4	10.4

a. Based on market expectations of the 1-week rate in late April 2012
b. See Box 1 for an explanation of the effective household borrowing rate.

The control scenario (**Table 1**) represents a stable macroeconomic environment in which the unemployment rate and the duration of the period of unemployment are unchanged throughout the simulation period. A gradual increase in the overnight rate is assumed, but the effective borrowing rate declines slightly as some fixed-rate mortgages are renewed at current rates (which are lower than the previous rates on the maturing debt of some households).

In the other scenarios, unemployment increases by 3 percentage points, and the average duration of unemployment rises by six weeks (similar to the assumptions used in past FSR stress-testing exercises, but in a multi-year context). Under the scenario with a short-lived unemployment shock, unemployment rises for only one year and then returns to control. In the persistent-shock scenarios (both with and without asset dynamics), unemployment increases gradually but remains elevated (**Table 1**). Under the persistent-shock scenario with asset dynamics, the prices of stocks and mutual funds decline from their starting point by a cumulative total of 28 per cent and 20 per cent, respectively.[13] Given that near-cash assets are largely unaffected, the total average cumulative effect of a change in asset prices in this scenario is a decline of about 15 per cent in household financial wealth.[14] Aggregate annual savings in all scenarios are kept at about 3 per cent of disposable

Chart 2: The average duration of periods of unemployment: Assumptions

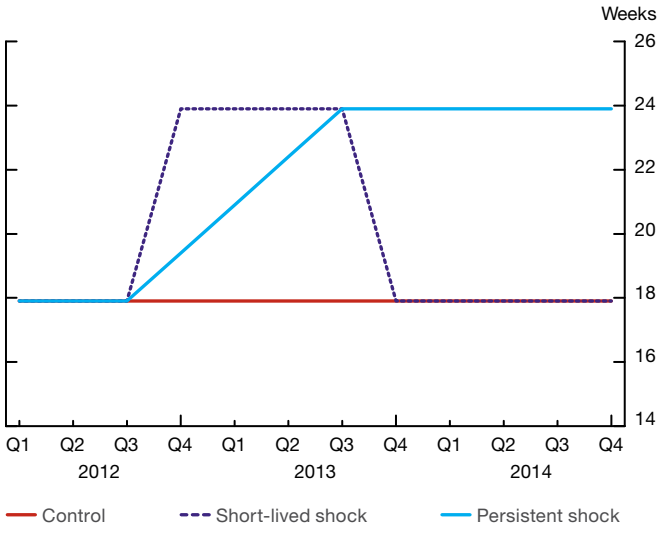

13 These figures reflect a market return that is comparable with the average of the five worst peak-to-trough declines in the Toronto Stock Exchange since the late 1980s.

14 The average return of -15 per cent on total financial assets is a weighted average of returns on individual classes of financial assets, with weights given by the aggregate holdings in household portfolios.

income.[15] The shock scenarios hold the policy rate constant to permit an assessment of the impact of these shocks on household vulnerability in the absence of mitigating policy actions. Nonetheless, the effective household borrowing rate increases in response to a rise in risk premiums of about 220 basis points.[16]

While HRAM does not explicitly model the spillover effects between the financial and real sectors, judgment can be used to incorporate this facet into the design of the scenario. For example, in the persistent-shock scenario shown in **Table 1**, weaker labour market conditions are accompanied by a tightening in lending standards that leads to higher risk premiums, a significant moderation in credit growth and lower house prices. These changes in financial conditions result in weaker growth in household spending and therefore amplify the decrease in aggregate income. These real-financial linkages can be implicitly captured in the scenario by choosing suitably severe assumptions for the unemployment rate and income growth.[17]

Simulation results

Under the control scenario, the share of vulnerable households (**Chart 3**), their share of debt (**Chart 4**) and arrears (**Chart 5**) are all broadly unchanged.

Increases in unemployment lead to greater vulnerabilities and arrears. Although the short-lived unemployment shock leads to a sharp rise in the two vulnerability measures and loans in arrears, most of these effects are temporary. These measures nonetheless return to a level that is somewhat higher than the control case as some households renew fixed-rate mortgages during the higher-rate period of the shock. For the two variants of the persistent-shock scenario, both measures of household vulnerability rise to a higher level, as do loans in arrears. The key difference in the results given by these scenarios is that loan arrears are lower when asset dynamics are allowed (**Chart 5**). This is explained in more detail below.

These results illustrate the key improvements in HRAM. Unlike the previous version, HRAM can now describe arrears over the entire simulation. The total effect of excluding asset dynamics from the shock scenario can be seen by comparing the green and blue lines in **Chart 5**. The green line shows the persistent-shock effect when there are no household savings from current income or returns on financial assets, while the blue line does account for these features. The net effect of asset dynamics is to mitigate the rise in arrears, indicating that the asset-price channel is dominated by the savings channel.

The relatively small impact of financial asset prices in the model can be explained by examining the balance sheets of households that go into arrears. These households typically start with low levels of assets that

Chart 3: Share of vulnerable households

Households with a debt-service ratio ≥ 40 per cent

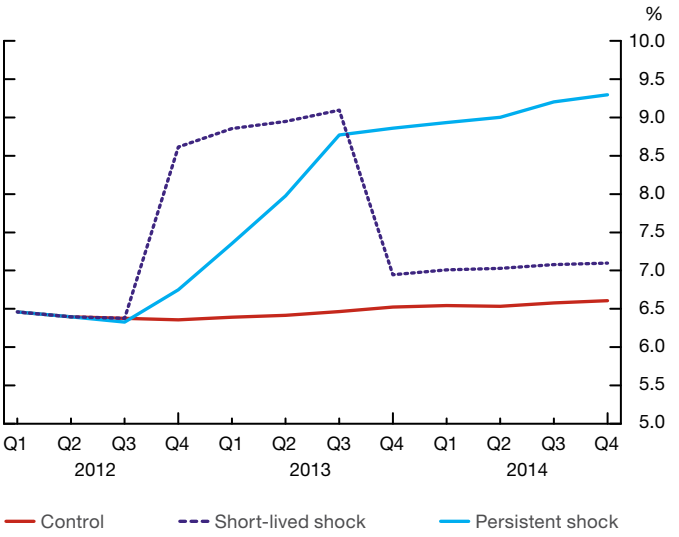

Source: Bank of Canada calculations

Chart 4: Share of debt held by vulnerable households

Percentage of total debt held by households with a debt-service ratio ≥ 40 per cent

Source: Bank of Canada calculations

15 Generally, the savings rate would increase as unemployment rises, as households become more precautionary. Keeping the savings rate unchanged excludes this mitigating effect.

16 The effective borrowing rate shown in Table 1 rises by less than the risk premium because only a fraction of fixed-rate debt is renewed in each period.

17 A longer-run objective for the development of HRAM is to model the real-financial linkages more formally.

Chart 5: Rate of household loans in arrears

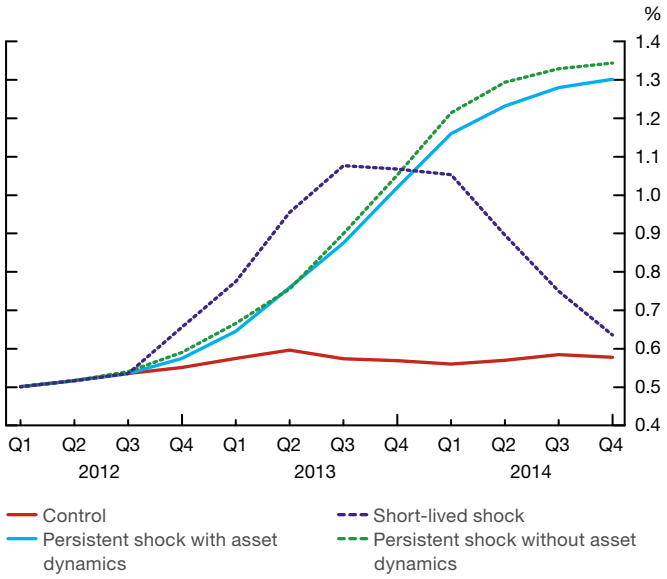

Source: Bank of Canada calculations

are generally in the form of less-risky investments and, therefore, are little affected by changes in equity prices. Conversely, households with high levels of risky assets also tend to have high levels of near-cash liquid financial assets (e.g., money market funds or savings accounts), which are largely unaffected by an asset-price shock. Even a significant shock is unlikely to push them into immediate distress.[18] Together, these facts show why movements in asset prices have a subdued impact on loan arrears in HRAM. On the other hand, relatively modest levels of savings can help to prevent financial distress for the most at-risk groups.

Sensitivity Analysis and Historical Evaluation

In this section, we outline two exercises that were conducted to assess whether HRAM yields reasonable results. First, we examine the sensitivity of the model to changes in the key assumptions for the scenarios. We then run a stylized scenario to evaluate the model's response to a historically extreme macroeconomic downturn.

Sensitivity analysis

For the sensitivity analysis, we modify the unemployment and interest rate assumptions without making any further changes to the other assumptions in the control scenario.[19] **Table 2** summarizes the impact of these changes on the rate of loan arrears by the third year of the simulation. The effect on arrears is somewhat greater than the proportional change in unemployment; for example, a 2-percentage-point increase in unemployment corresponds to an 85 per cent increase in arrears, relative to the most recent historical observation of 0.5 per cent in 2011Q4. The multi-year unemployment shock leads to a cumulative financial strain on households, contributing to the greater-than-proportional effect. Overall, the exercise confirms the significance of unemployment in driving arrears.

The impact of a given increase in interest rates becomes more pronounced when it is combined with higher unemployment. In the most extreme example shown in **Table 2**—where unemployment rises by 6 percentage points and household borrowing rates increase by 400 basis points—the rate of arrears almost quadruples, to a level of about 2.0 per cent (a 290 per cent increase from the starting point of 0.5 per cent).[20] While credit growth would, in reality, slow down in response to the change in interest rates, we would still expect the increase in arrears to be significant, given the severity of the scenario.

Table 2: Percentage increase in arrears as a result of changes in unemployment and interest rates[a]

Unemployment (percentage points)	Interest rates (basis points)		
	+0	+200	+400
+0	15	34	50
+2	85	111	141
+4	136	176	215
+6	191	221	290

a. Measured as the average effect in the third year, relative to the starting-point arrears of 0.5 per cent, with other assumptions taken from the control scenario

The model's response to a 1980s-style recession

We use the recession that occurred in the early 1980s as the basis for our historical model-evaluation scenario. To test the model, we specify a shock that is similar to

18 Despite this observation, the longer-term financial well-being of Canadians could still be significantly affected by such a shock.

19 Assumptions are modified in parallel shifts relative to the control scenario. For example, if unemployment is increased by 2 percentage points (all other assumptions held at control scenario), then the unemployment rate would be higher by this amount (relative to control) for each period of the simulation.

20 We expect that this figure should be well below the 7 per cent to 8 per cent rate of arrears seen in the United States since 2009. Before the crisis, from 1999 to 2006, loans in the United States that were 90 days or more in arrears represented, on average, 2.2 per cent of the total loan balance, more than five times higher than comparable figures for Canada. Stronger provisions for lender recourse in Canada contribute to this difference. In addition, the distribution of the financial profiles of homeowners in the United States had more pronounced vulnerabilities, owing to the expansion in subprime lending.

that episode, and then examine whether the simulated rate of loan arrears is broadly comparable with the level observed then (i.e., a peak of about 1 per cent). The scenario includes an increase in unemployment of 6 percentage points, and household borrowing rates that peak at 21.0 per cent. The impact of a given shock on arrears will vary over time, depending on the underlying vulnerabilities at the time of the shock. Since the current distributions across households of key determinants of vulnerability (e.g., asset levels and the DSR) are much different from those observed in the 1980s, we make adjustments so that these conditions are broadly representative of the earlier period. For example, we rescale debt levels to reflect the fact that the aggregate debt-to-income ratio was much lower (close to 80 per cent) in the early 1980s.

In this scenario, the model predicts that loan arrears would peak at about 1.2 per cent. Available data on arrears for the early 1980s—which are restricted to uninsured mortgages—show a peak of just above 1.0 per cent. Since the peak would likely have been somewhat higher if the historical data included consumer debt and insured mortgages, the model is able to broadly generate the level of arrears experienced during this recession.

The results highlight the importance of starting-point distributions for the DSR and liquid financial assets. Using the current distributions without any adjustments, a similar 6-percentage-point increase in unemployment, as well as higher interest rates, would lead to a significantly greater increase in arrears (**Table 2**).

Conclusion

Stress tests using microdata are a significant component in the assessment of the financial stability risk related to household balance sheets. While aggregate measures can describe important sectoral trends, it is at the micro level that we can better assess the potential change in loan arrears under an adverse-shock scenario.

This report highlights the methodological advances made by the Bank with respect to its stress-testing framework for household financial stability. These include adding the capacity to simulate a multi-year integrated unemployment shock and allowing household assets to evolve over the simulation. The sample scenarios used in this report are illustrative only; the revised stress-testing framework has the flexibility to consider a broader array of alternative assumptions, allowing for an improved assessment of household vulnerabilities.

Further extensions to HRAM are planned to strengthen the empirical foundation of the scenario design. A key example is to base the assumed path for credit growth more closely on empirical findings on the relationship between credit and such variables as interest rates, income and house prices, rather than on judgment. As well, the method of implementing the unemployment shock will be revisited. Currently, unemployment is distributed randomly across all employed households without incorporating household-specific factors into the distribution. Finally, the behavioural detail in the model could be further extended. For example, the simulation does not currently allow distressed households to sell their houses. If they are allowed to do so, they can potentially avoid default, leading to lower arrears. This extension would need to consider how the feasibility of this option would change in a severe housing-market downturn with falling house prices and slower market turnover.

References

Djoudad, R. 2010. "The Bank of Canada's Analytic Framework for Assessing the Vulnerability of the Household Sector." Bank of Canada *Financial System Review* (June): 57–62.

Djoudad, R. 2012. "A Framework to Assess Vulnerabilities Arising from Household Indebtedness Using Microdata." Bank of Canada Discussion Paper No. 2012-3.