Are there significant disparities in debt burden across Canadian households? An examination of the distribution of the debt service ratio using micro-data

Umar Faruqui[1,2]

1. Introduction and summary

The household debt-to-income ratio in Canada has increased from 110 per cent in 1999 to 123 per cent in 2005. This rapid debt accumulation has raised concerns about the ability of households to deal with debt payments if interest rates increase or if they face a negative economic shock.

The debt service ratio (DSR) is one metric to gauge the burden of debt servicing for households and has been the focus of increased scrutiny as of late. In Canada - up until now - most of the analysis of the household DSR has been based on aggregate data, which show that households' debt servicing burden is near historic lows. However, these aggregate data average across all households and can mask information about the distribution of the debt burden.

A number of studies have examined household indebtedness using micro data. These include Canner et al. (1995) and Barnes and Young (2003) for the U.S., May, Tudela and Young (2004) for the U.K. and Herrala (2006) for Finland. The studies suggest a number of common results. In particular, they find that: (a) the micro data evidence on household financial health matches up relatively well with the information from aggregate data, and (b) households that have high debt levels are also those that can most safely bear that debt burden.

In this paper we examine the distributional properties of the DSR for Canadian households using micro data, which show that:

- Debt and asset holdings of households are relatively well-matched.
- The incidence of variable rate debt has increased over the last seven years, with older and richer households holding a larger proportion of their debt at variable rates than younger, poorer households.
- The distribution of the DSR has not changed much since 1999.
- The density of households in the vulnerable tail of the DSR distribution has decreased since 1999, especially for lower-income households.

Overall, we find that the micro data support inferences based on the aggregate data: despite the increase in the debt-to-income ratio since the late 1990s, households' financial health remains sound.

---

[1] The views expressed in this paper are those of the author. No responsibility for them should be attributed to the Bank of Canada.

[2] Monetary and Financial Analysis Department, Bank of Canada, ufaruqui@bankofcanada.ca. I thank Simon Lai for his excellent research assistance, and Allan Crawford, Pierre Dugay, Ben Fung, Dinah Maclean and Virginie Traclet for their comments and suggestions. All remaining errors and omissions are my own.

The remainder of this paper is organized as follows. The next section provides some background on the DSR. Section 3 describes the micro data used for the analysis. Section 4 outlines the recent update to the Bank of Canada's aggregate DSR. Section 5 examines the distribution of the DSR across Canadian households, and section 6 concludes with a discussion of the results and a list of future steps for this topic.

2. Background

2.1 What is the DSR?

The DSR measures the proportion of disposable income (net of taxes and transfers) that households must devote to servicing their debt obligations. The ratio can be calculated two different ways. The traditional DSR calculation only considers interest payments on debt as the cost of debt for households. An alternate measure includes both interest payments as well as principal repayments in debt servicing costs. The Bank of Canada's (henceforth Bank) aggregate DSR measure (discussed below) adheres to the more traditional definition and considers only the interest service burden of debt for the household sector.

Both definitions of the DSR have their advantages and disadvantages. The interest-only DSR (IO-DSR) is often easier to calculate from the data as it requires less detailed information about debt repayments. Changes in the IO-DSR are also simpler to interpret and can be used to isolate the impact of interest rate changes on the household's debt burden. However, the IO-DSR captures only a portion of the cost of debt, as principal repayments can form an important component of debt obligations, especially for mortgage debt. The IO-DSR may, therefore, provide a misleading picture of the household debt burden in a high debt, low interest rate environment. Under these conditions, the IO-DSR may understate the actual burden of servicing debt on households. The DSR measure that includes interest and principal repayments (IP-DSR) is an arguably better measure of the household debt burden as it includes all debt-related payments that a household has to make. The IP-DSR is, however, often harder to estimate since it requires more detailed information on non-revolving loans.

The IP-DSR is our preferred measure of households' debt service burden and we use this measure for the analysis of the distribution of the DSR in section 5. However, when discussing the Bank's aggregate DSR, we are obliged to work with the IO-DSR due to data limitations.[3]

2.2 Why is the DSR of interest to central banks?

The DSR provides information for both monetary policy and financial stability. For monetary policy, the DSR can be used to estimate the proportion of household disposable income available for discretionary purchases. When the ratio is high (relative to some benchmark), households have fewer funds to spend on current consumption. Furthermore, households with a high debt service burden are more likely to be adversely affected by a negative shock, such as an employment or life event shock. If there are a large number of households with a high DSR, current period consumption may be more adversely affected by a negative shock than otherwise. Finally, a high debt service burden may constrain a household's access to credit affecting its ability to smooth consumption over time.

For financial stability the DSR can be used to measure the household sector's ability to service its debt over time. If this ability deteriorates - i.e. the DSR rises significantly following

[3] Work is underway to expand this aggregate measure to include principal repayments as well as interest payments.

a negative shock - financial institutions may face rising loan arrears and/or personal bankruptcies. This would translate into a deterioration in their asset positions, lowering their profitability and potentially making the financial system more vulnerable. For these reasons, the DSR is closely followed by central banks.

While the aggregate DSR can provide useful information about the debt service burden on the average household, it provides no insight into the distribution of that debt burden across households. An analysis of the DSR distribution requires micro household data and is a useful complement to the aggregate measure. If the DSR distribution has a 'fat' right tail, it means that a large proportion of households have a high DSR. Under these circumstances a negative macro-economic event would probably lead to a larger impact on aggregate consumption of households than if the DSR distribution was not skewed. Furthermore, there would be a higher risk to financial stability as a larger number of households may be at risk of default or bankruptcy than otherwise.

2.3 Calculation and current use of the DSR at the Bank

Up until now, the Bank has focused on the IO-DSR based on aggregate data. The reliance on aggregate data has been due, in large part, to the paucity of comprehensive and timely micro data on household balance sheets.[4] The aggregate DSR has been used by the Bank to assess the implications of rising interest rates on the household debt servicing burden via scenario and stress-testing analyses.[5]

The Bank's IO-DSR is based on aggregate debt and disposable income data from Statistics Canada.[6] There are two steps in calculating the ratio from these data. First, the aggregate debt data are broken down into more recognizable loan categories and second, an effective interest rate is applied to each loan category in order to estimate the interest servicing cost of the debt. These steps require assumptions about the relative proportion of each loan category in total outstanding debt and interest rates paid on each type of loan. The assumptions prior to 1999 are based on a mix of anecdotal and survey evidence, while the post-1999 assumptions are based on information from the Canadian Financial Monitor survey (discussed in Section 3). The resulting DSR is denoted by the following equation:

$$IO-DSR_t = \frac{\sum_i (r_{i,t} \times Debt_{i,t})}{DI_t} \quad (1)$$

Where:

'*i*' = different categories of household loans,

'*r*' = effective interest rate on each category of loan,

'*Debt*' = outstanding balance on each category of loan, and

'*DI*' = aggregate disposable income for the household sector.

---

4 For example, the Survey of Financial Security - a survey of household finances compiled by Statistics Canada - is only available for selected years (e.g. 1984, 1999 and 2005) and lacks detailed information on the debt side of household balance sheets to construct a DSR measure.

5 See December 2004, *Financial System Review* for more details.

6 Aggregate debt data are from the National Balance Sheet Accounts, while disposable income data are from the National Income and Expenditure Accounts.

3. The data

Data from the Canadian Financial Monitor (CFM) forms the basis for the analysis presented in this paper. These data not only allow us to refine the Bank's aggregate IO-DSR but also let us explore the distributional aspects of the DSR (based on the IP-DSR measure).[7]

This section describes the CFM data and presents selected stylized findings about household asset and debt distributions using these data.

3.1 Description of the CFM survey

CFM is a household survey conducted by Ipsos Reid Canada and provides detailed balance sheet information. The survey, which started in 1999, has a sample size of approximately 12,000 households annually who respond through a mail-in form.[8,9] Currently we have seven years of survey data, from 1999 to 2005.

The survey content has remained roughly unchanged since the inception of the survey in 1999. The 2005 survey consisted of ten sections of questions: five sections on assets, three on debt and one each on banking behaviour and household characteristics. The household characteristics section collects information on the age group of the household head, family income, family size and marital status of the household head, amongst other things. Up until now, CFM data have been primarily used by Canadian financial institutions for market research.

3.2 Comparing CFM data with other datasets

CFM data compare favourably with other Canadian household survey data. CFM has one year of overlap (1999) with Statistics Canada's Survey of Financial Security (SFS)[10] and for that year the two surveys match up relatively well, especially on the debt side of household balance sheets.[11]

---

7 While data limitations only allow estimation of the IO-DSR at the aggregate level, CFM allows the estimation of both the IO- and IP-DSRs. We choose to focus on the IP-DSR measure from CFM for the distributional analysis for the reasons outlined in section 2.1.

8 The survey has a monthly distribution target of 1,600 in January, February and March and 800 in each of the remaining months of the year. Respondents are given incentives for completing and returning valid surveys including draws for prizes. The response rate for the survey was roughly 35 percent for the 2003-2005 period.

9 The survey has both a cross-sectional and panel dimension. Of the approximately 12,000 households included in the sample in each year about half are from a rotating panel. For the 1999-2005 period, there are approximately 56,000 households for whom we have more than one observation and 3,000 for which we have a full time series of seven observations. Future work with CFM will examine the usefulness of this panel aspect of the data.

10 While the periodic nature of the SFS does not make it suitable for ongoing analysis, it is nonetheless a useful quality check for CFM. In fall 2006, the 2005 SFS survey data will become available allowing a more comprehensive robustness check of the CFM data.

11 Based on internal analysis by Geoff Wright and Nicholas Brewer of the Bank of Canada.

Table 1

Annual growth in debt - macro vs. micro data

	Aggregate data						Micro data		
	Statistics Canada (NBSA)			Bank of Canada			CFM		
	Consumer	Mortgage	Total	Consumer	Mortgage	Total	Consumer	Mortgage	Total
2000	9.2%	4.3%	6.1%	12.6%	4.8%	7.1%	11.8%	–0.1%	2.9%
2001	8.8%	4.4%	6.1%	6.8%	4.0%	4.9%	7.0%	7.8%	7.6%
2002	7.2%	5.5%	6.2%	6.5%	7.4%	7.1%	8.1%	7.3%	7.5%
2003	7.1%	6.5%	6.7%	8.7%	8.2%	8.3%	–0.2%	4.1%	2.9%
2004	7.0%	6.7%	6.8%	9.9%	9.8%	9.8%	10.8%	–1.2%	2.0%
2005	7.2%	6.9%	7.0%	12.2%	10.1%	10.8%	8.7%	9.8%	9.5%
Avg: 2000-05	**7.7%**	**5.7%**	**6.5%**	**9.5%**	**7.4%**	**8.0%**	**7.7%**	**4.6%**	**5.4%**

There are two different sources of aggregate household debt information for Canada: Statistics Canada's National Balance Sheet Accounts (NBSA) and the Bank of Canada's credit data based on bank returns.[12,13] The estimated debt levels from CFM are noticeably lower than both sets of aggregate data.[14] While there is more noise in the micro data than in the aggregate data, the growth in household debt from CFM and the aggregate data sources show a roughly similar trend of increasingly rapid household debt accumulation over the 1999-2005 period (Table 1).[15]

3.3 Caveats about CFM

While the debt-side of the CFM data compares relatively well with other micro data and the aggregate data, the survey is not without its weaknesses. First, we find that the asset side information from CFM is noticeably different from the SFS (for 1999).[16] Second, CFM

[12] Bank credit data are available as part of the *Weekly Financial Statistics* publication at http://www.bankofcanada.ca/en/wfsgen.html.

[13] There are two primary differences between the two sets of aggregate data. The first difference involves the definition of the household sector: the NBSA data includes unincorporated businesses into the household sector, while the Bank's data does not. Second, the classification of debt is different across the two sets of aggregate data: NBSA classifies debt by type of borrower, while the Bank's data classifies according to the use of funds. For example, if a business borrows money to build an apartment building, it would be included in household debt under the Bank definition but not under the NBSA guidelines.

[14] Aggregate debt levels implied by the CFM micro dataset are generally about 80 per cent of those from the aggregate sources. This could be a result of two factors. First, there are conceptual differences between the micro and macro data. In particular, the NBSA's definition of the household includes unincorporated businesses and the Bank of Canada includes loans by businesses for residential investment. CFM only includes borrowing by households. Second, evidence from other studies suggests that households tend to under-report both their debt and asset holdings in surveys.

[15] A comparison of the mean IO-DSR from CFM vs. that from the aggregate data is presented in Appendix 1. We find that the estimates from CFM and aggregate data show very similar trends over the 1999-2005 period.

[16] While it is not clear what factors account for the differences between total asset levels from CFM and SFS, we suspect that survey collection method may be a factor.

information on loan, asset balances and household income are collected as ranges and not as a point estimate, which may introduce noise in information extracted from CFM. Using the mid-point of the range as a point estimate of loan, asset balances and income is a second-best solution, especially when the ranges become wide. Third, the income groups are top and bottom coded, which means, for example, that any household with an income greater than $150,000 is coded as being in the ">$150,000" income group. Similar problems arise with loan and asset balances. Finally, since CFM is a mail-in survey, it is likely to contain more internal errors and inconsistencies than, for example, phone surveys where the interviewer can probe the household to ensure that they answer in an internally consistent manner.

3.4 Selected stylized findings from CFM

For the presentation of the stylized findings we define the following groups according to household characteristics: (1) income groups: low income (gross family income of less than $35,000), middle income ($35,000<= income <$70,000), and high income (income>=$70,000), and (2) age groups: young (age of household head less than 35 years), middle-aged (35<= age <50), and old (age>=50).[17] Table 2 shows the distribution of Canadian households by income and age groups.

Table 2

Proportion of population in income and age group[1]

	Middle			**Sum**
	Young	**Aged**	**Old**	
Low-income	0.08	0.10	0.19	**0.37**
Middle-income	0.09	0.13	0.13	**0.34**
High-income	0.06	0.14	0.09	**0.29**
Sum	**0.23**	**0.36**	**0.41**	

[1] Pooled data, 1999-2005.

Most Canadian households carry some form of debt. The proportion of households with positive debt levels has declined slightly from 77 per cent in 1999 to around 75 per cent in 2005.[18] Table 3 shows the distribution of debt and assets by age and income groups. These results indicate that debt holdings differ markedly by demographic and financial characteristics, and are broadly consistent with predictions from the life-cycle theory of consumer behaviour. In particular, the table shows that: (1) middle-aged households hold the majority of total outstanding debt, even though they form a smaller portion of the population than older households, (2) debt holdings initially increase with the age of the household but then decline as the households reach old age, and (3) richer households hold a relatively

[17] The definition of income groups is arbitrary but is broadly consistent with definitions used by Statistics Canada (Statistics Canada, "*Income in Canada, 2004*". Cat. # 75-202-XIE). According to Statistics Canada, households earning below 50 per cent of the median income are considered low income. The median gross income for Canadian households was $63,100 in 2004.

[18] This is comparable to figures for the U.S., which show that roughly 74 per cent of U.S. households hold some form of debt (Barnes and Young, 2003).

large proportion of debt given their size in the population. On the other side of the balance sheet, asset holdings increase with both income and age. A number of other studies (e.g. Edelberg and Fisher (1997) and Reserve Bank of Australia (2003)) have found similar results for other developed countries.

Table 3

Proportion of debt and assets held by household groups[1]

	Middle			Sum
	Young	**Aged**	**Old**	
Debt				
Low-income	0.04	0.04	0.04	**0.11**
Middle-income	0.12	0.15	0.07	**0.34**
High-income	0.15	0.29	0.11	**0.54**
Sum	**0.31**	**0.48**	**0.21**	
Assets				
Low-income	0.01	0.03	0.10	**0.14**
Middle-income	0.05	0.10	0.18	**0.32**
High-income	0.07	0.23	0.24	**0.54**
Sum	**0.13**	**0.36**	**0.52**	

[1] Pooled data, 1999-2005.

Households with high debt relative to assets can be more vulnerable to shocks[19] and those households with a high proportion of debt at variable rates would have debt service payments, which are very sensitive to any increase in interest rates.[20] Therefore, two questions of particular interest for policy makers that CFM can help address are: (1) Do high-debt households also have high asset balances? and (2) what type of debt (variable vs. fixed, for example) is held by different categories of households?

Figure 1 shows the distribution of the debt-to-asset ratio (DAR) for households with non-zero debt. While the DAR from CFM should be interpreted with some caution for the reasons outlined in section 3.3, the trends in the data can provide useful information. The strong left skew in the DAR distribution suggests that (for the most part) debt and asset holdings of households are well matched.

[19] Debt-to-liquid assets may be a more revealing ratio for some purposes. We leave that analysis for future work.

[20] For certain types of variable rate debt (e.g. variable rate mortgages) the monthly payments don't change when interest rates change. Instead, the proportion of the fixed monthly payment that goes towards repayment of principal adjusts accordingly. Our analysis abstracts from this point.

Figure 1

Distribution of debt-to-asset ratio across households[1]

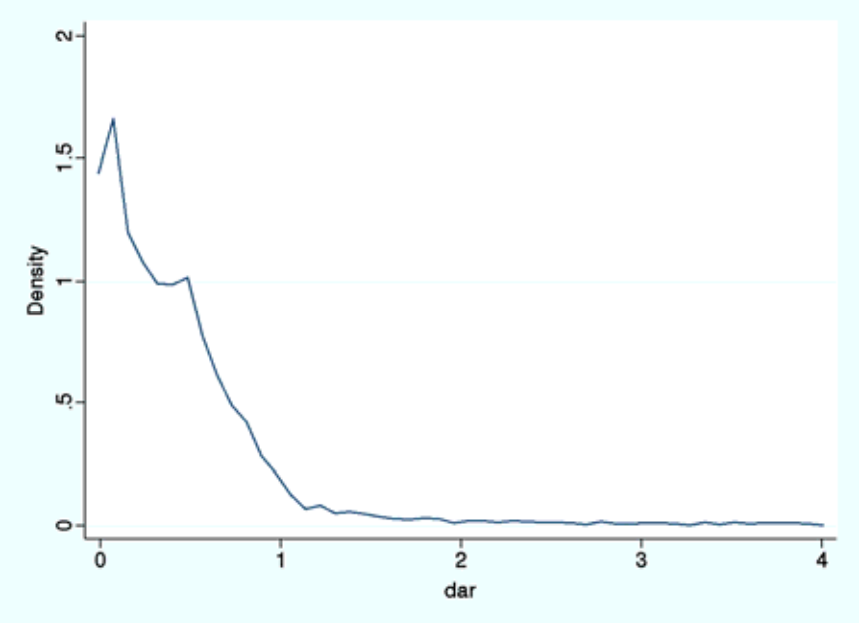

[1] Pooled data, 1999-2005. Excludes households with zero debt and extreme outliers.

Closer examination of the tails of the DAR distribution shows that while the median of the distribution has declined since 1999, the density of households with very high DAR[21] has increased somewhat (Table 4).

Table 4

Selected DAR statistics by year[1]

	Median	Density: DAR >2
1999	0.32	4.7%
2000	0.33	5.1%
2001	0.33	5.2%
2002	0.36	4.6%
2003	0.31	5.2%
2004	0.30	6.6%
2005	0.29	6.8%

[1] Excludes households with zero debt and extreme outliers.

[21] We define vulnerable households with regard to the DAR as households with a DAR above two, which is consistent with the fact that, from 1987 to 2004, the average DAR of insolvent households in Canada has hovered around two.

An examination of the breakdown of debt by type of interest rate for different age and income groups shows that older and high-income households are more likely to carry variable-rate debt than other households (Table 5).

Table 5

Proportion of debt on variable rates[1]

Low-income	0.16	0.17	0.24
Middle-income	0.15	0.19	0.27
High-income	0.18	0.27	0.34

[1] Pooled data, 1999-2005. Excludes households with zero debt.

We also find that (confirming anecdotal information to this effect) the incidence of variable-rate debt has increased from 14 per cent in 1999 to 32 per cent in 2005.[22] The data show that the increase in the holding of variable-rate debt was widespread across income and age groups. This increase in the popularity of variable rate debt can be accounted for, in part, by the rising spread between long and short-term interest rates in Canada over the 1999-2005 period.[23]

Summary

Overall, the stylized facts from CFM suggest that:

- The debt and asset positions of households are reasonably well matched, with richer households holding a large proportion of overall household debt.
- Older households hold a larger share of their debt at variable rates than younger households.
- The share of variable debt holdings by all household groups has been increasing over recent years. It could be argued that the higher incidence of variable rate debt makes households' financial obligations more sensitive to interest rate increases. However, as the most recent anecdotal information from our contacts at financial institutions suggests, households do adjust their borrowing behaviour (and relatively quickly) in response to movements in interest rates. Since September 2005, the short-term rate in Canada (as proxied by the prime rate) has increased by 175 basis points and the spread between long and short-term interest rates has narrowed. At the same time, households have swiftly moved away from variable-rate and towards fixed-term debt.[24]

[22] Variable rate debt includes variable rate mortgages, leases and other consumer loans on variable rates. Fixed rate debt includes credit card debt, fixed-rate mortgages and consumer loans on fixed rates.

[23] The interest rate on variable-rate debt products is usually based on the short-term interest rate, while rates on fixed-rate debt products are based on the relevant long rate.

[24] This is especially true for mortgage debt where it is relatively easy for households to switch from variable to fixed rates.

4. Update to the aggregate IO-DSR

Our aggregate IO-DSR for the post-1999 period has been recently updated using information extracted from CFM. The new aggregate IO-DSR estimate incorporates two main improvements and updates to the underlying assumptions. First, consumer debt is now partitioned into finer categories: credit card debt, secured personal lines of credit (PLCs), unsecured PLCs, other personal loans (fixed rate), other personal loans (variable rate), and automobile leases.[25] The finer breakdown allows us to better capture the impact of important changes in consumer credit, notably the substitution away from personal loans and credit card debt to secured PLCs, since the late 1990s.[26] Second, the assumptions regarding the breakdown of mortgage debt and effective interest rates on different types of mortgage products are updated. This update better reflects the impact of the shift from variable to fixed rate mortgages and the presence of prevalent discounting on variable rate mortgages (as seen over recent years) on the IO-DSR.

Figure 2

Aggregate IO-DSR for the Canadian household sector

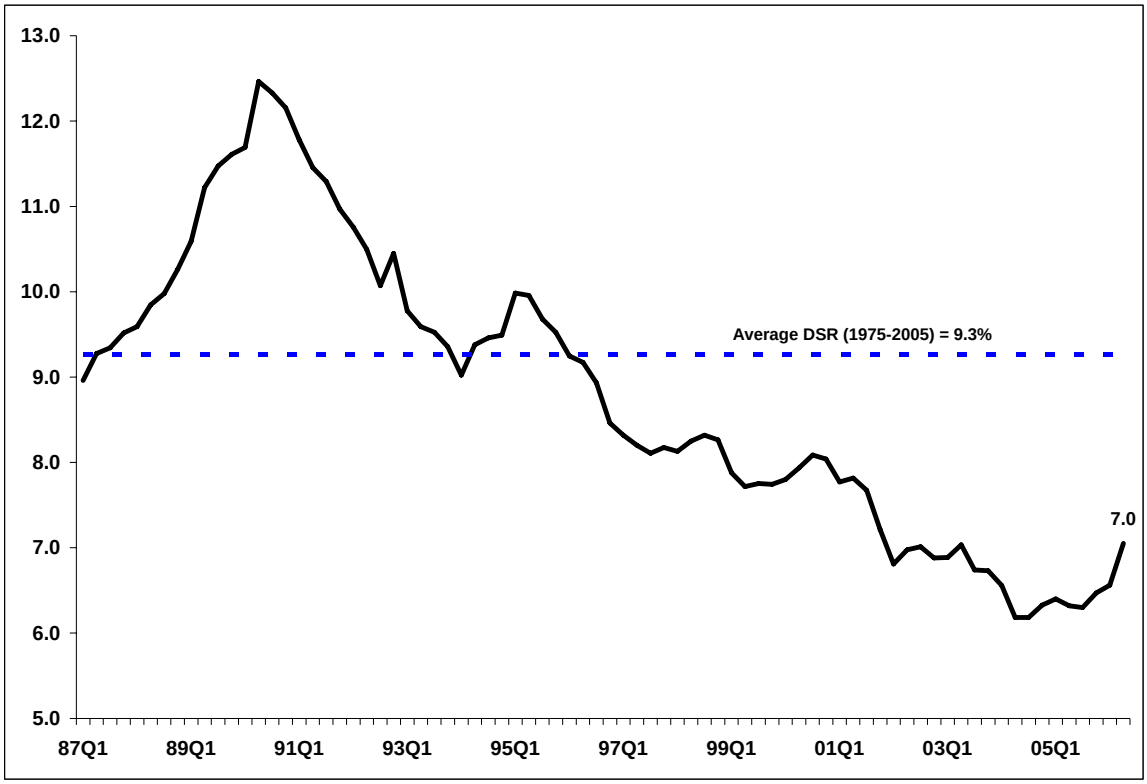

According to the aggregate estimate, the burden of servicing existing debt (interest-only) for Canadian households was well below the historical average in 2006Q2, suggesting that the household sector is in good financial health. Since 2000, the IO-DSR has trended down even as the debt-to-income ratio has increased rapidly. These two observations are reconciled by the decline in the effective interest rate on debt over this period.

---

[25] Under the previous DSR assumptions, consumer debt was only subdivided into two categories: variable and fixed rate loans.

[26] This shift is important because secured PLCs - which bear lower interest rates compared to other forms of consumer credit - have contributed to reducing households' debt-servicing costs.

While the aggregate IO-DSR is a useful indicator of household financial well-being, it has some shortcomings. These include:

- Arguably what matters for monetary policy and financial stability is the total debt burden on households including principal repayments. Future work at the Bank will try to extend the aggregate debt burden measure to include principal repayments but at the moment there is a lack of adequate information for this calculation.
- The aggregate DSR measure can mask potential issues with the distribution of the debt service burden across households. For example, if the distribution of the DSR was skewed to the right, it would mean that there is a high incidence of households with an elevated debt service burden. This, in turn, may have implications for both monetary policy and financial stability as it may affect the response of the household sector to macro-economic shocks. Since the aggregate DSR estimate provides no information about the underlying distribution of that debt burden across households, it is important to augment the aggregate analysis with micro data analysis of household indebtedness. An analysis of the DSR distribution across households is provided in the next section.

5. Distributional analysis of the household debt burden

In this section we address three main questions using CFM data: (1) how is the debt service burden distributed across households? (2) has this distribution of the DSR changed over recent years? and (3) what is the density of households in the vulnerable tail of the DSR distribution?

As mentioned above, the total debt burden (i.e IP-DSR) is arguably a more useful indicator than the interest-only measure (IO-DSR). Due to data constraints on the aggregate level, calculating an aggregate IP-DSR is quite hard. By contrast, CFM has data on total debt payments including principal repayments thus allowing a calculation of the IP-DSR. We use the IP-DSR for the distributional analysis presented in this section.[27] The IP-DSR for each household ('*j*') and year ('*t*') is estimated as follows from the micro-data:[28]

$$IP-DSR_{t,j} = \frac{\sum_{i} payment_{i,j,t}}{GI_{t,j}} \qquad (2)$$

Where:

'*i*' = mortgage loans, personal lines of credit, auto loans, outstanding credit card balance, other personal loans,

'*j*' = household ID,

'*payment*' = estimated annual payment to service loan,[29] and

'*GI*' = gross household income.

---

[27] As shown in Appendix 2, both the IO- and IP-DSRs estimated from CFM show similar trends over the 1999-2005 period. Therefore, the general conclusions from our analysis of the IP-DSR should also hold for the IO-DSR.

[28] Households that provide incomplete information about loans (needed for the estimation) are excluded from the calculations. This filtering decreases the coverage to around 80% of households with non-zero debt for IO-DSR and 82% of eligible households for the IP-DSR calculation.

[29] The survey collects data on the most recent payment on each loan and the frequency of the payment. We estimate annual loan service payments by assuming that the last loan service payment is representative of the periodic payments and then annualizing the periodic payment using the given frequency of payment.

An important point to note in the formula above is that the CFM-based measure for IP-DSR uses gross income rather than disposable income in the denominator. This is because there isn't enough information in CFM to estimate disposable income of the household.[30]

5.1 DSR distribution across Canadian households

Figure 3 shows the distribution of the IP-DSR across all households, conditional on the household having a positive debt balance.[31] The figure shows while the distribution is positively skewed, the long right-hand tail is quite thin, i.e. a relatively small number of households have a high IP-DSR.

Figure 3

Distribution of the IP-DSR[1]

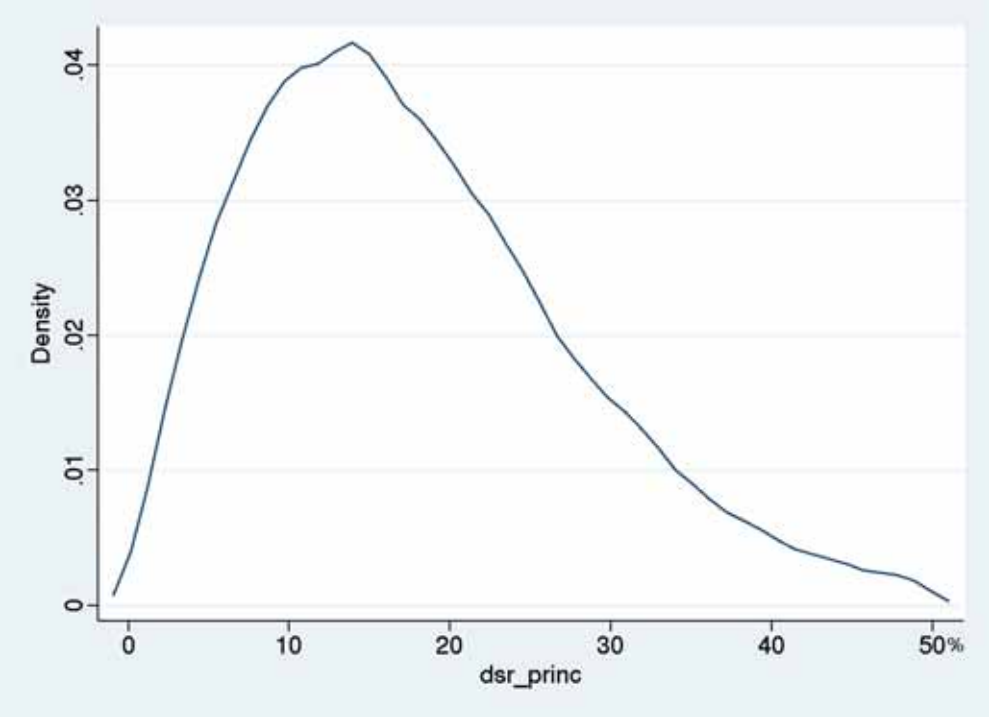

[1] Pooled data, 1999-2005. Excludes extreme outliers and households with no debt. Kernel density is the fitted density curve based on CFM DSR distribution data.

[30] The qualitative results from our analysis would be unlikely to change if we were able to use disposable income rather than gross income.

[31] This condition excludes those households with a zero DSR.

Figure 4

Distribution of the IP-DSR by income groups[1]

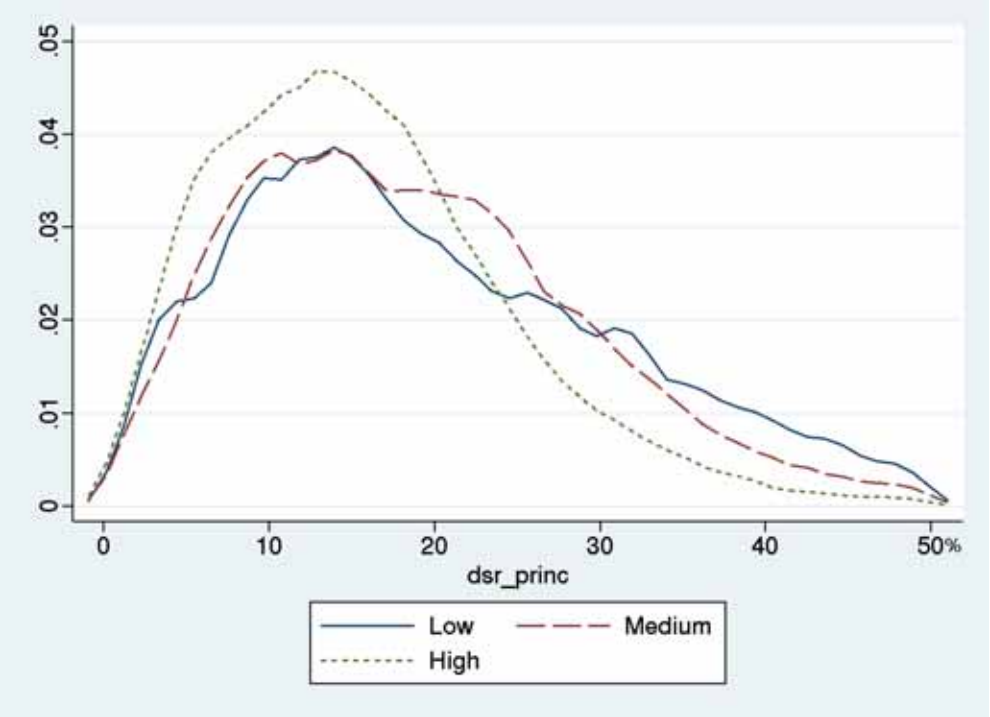

[1] Pooled data, 1999-2005. Excludes extreme outliers and households with no debt.

A priori, we would think that the distribution of the debt service burden would be different across households with different income. Indeed, the micro-data show that there is marked variation in the distributions of the IP-DSR for the three income groups. In particular, the lower income households have a more positively skewed distribution than higher income households (Figure 4).

Figure 5 plots the distributions of the IP-DSR (all households with positive debt balance) for 1999, 2002 and 2005. From this graphic, it is clear that the shape of the distribution has remained largely unchanged since 1999.[32]

An analysis of the higher moments of the distributions confirms this observation: the variance, kurtosis, and skewness of the distribution are very similar across the years (Table 6).

[32] The same trend is depicted if we examine the distributions by income groups.

Figure 5

Distribution of the IP-DSR, selected years[1]

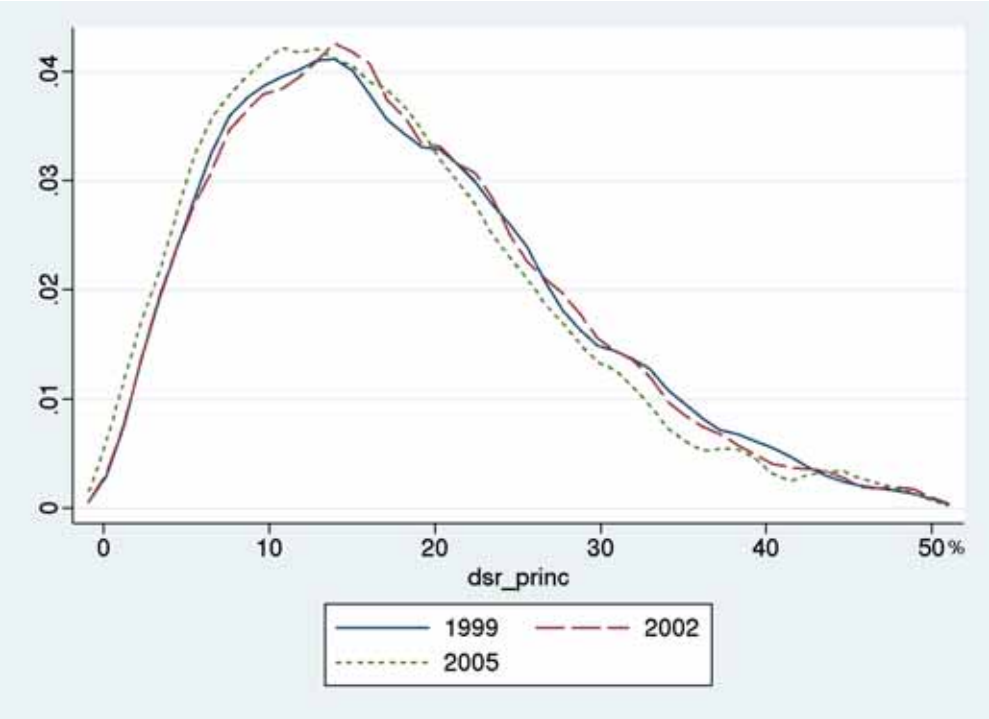

[1] Excludes extreme outliers and households with no debt.

Table 6

Moments of DSR distribution (by year)[1]

	Mean	Median	Variance	Skewness	Kurtosis
1999	17.77	16.27	98.96	0.65	2.94
2000	18.41	16.80	109.45	0.65	2.90
2001	18.33	16.80	104.52	0.63	2.92
2002	17.72	16.32	96.73	0.66	3.07
2003	17.85	16.44	99.16	0.67	3.05
2004	17.11	15.38	98.91	0.79	3.36
2005	16.69	15.28	94.89	0.76	3.32

[1] Moments of the conditional distribution (debt > 0 and excluding extreme outliers) of the IP-DSR.

However, while the higher moments of the distribution have not changed much over the last six years, the mean of the IP-DSR has decreased. Lateral shifts are important to our analysis as they influence the density of households in the 'vulnerable' tails of the distribution.

5.2 Density of households in "vulnerable" tail of the DSR distribution

One reason that the distribution of the DSR is of interest to policy makers is that it provides information on the proportion of households that are in a high-risk situation, i.e. households that are (relatively more) vulnerable to economic and other types of shocks. While there is no

universally accepted definition of the threshold for vulnerable tails of the DSR distribution, two commonly used thresholds are DSRs of 30 and 40 per cent.

The first threshold value is based on work by DeVaney (1994) who uses U.S. household data and probit analysis to show that having an IP-DSR greater than 30 per cent is an important determinant of future insolvency for a household. The second threshold (IP-DSR of 40 per cent) is based on anecdotal information from our Canadian bank contacts. Financial institutions often use the 40 per cent threshold to determine whether or not to extend credit to borrowers. One issue with the first threshold (IP-DSR of 30 per cent) is that it is expressed as a ratio to disposable income, whereas our IP-DSR measure from CFM uses gross income.[33] Since disposable income (on average) is about 75 per cent of gross income,[34] we can scale this threshold accordingly: the 30 per cent threshold is transformed into 23 per cent. In our analysis we use the scaled value of the thresholds and show the results using both the 23 and 40 percent thresholds as a sensitivity check on the results.

An examination of the density of households in the vulnerable tail of the IP-DSR distribution shows that this number has either remained roughly unchanged (IP-DSR 40 per cent) or decreased (IP-DSR 23 per cent) since 1999 (Table 7). These results are not entirely surprising as the shape of the DSR distribution has not changed since 1999 while the distribution has shifted to the left (i.e. the mean has decreased).

Table 7

Density in vulnerable tail (by year)[1]

	Households with debt >0		**All households**	
	DSR > 40%	**DSR > 23%**	**DSR > 40%**	**DSR > 23%**
1999	2.6%	29.3%	1.6%	19.3%
2000	4.6%	31.2%	2.9%	21.0%
2001	3.9%	30.6%	2.6%	20.7%
2002	3.0%	28.4%	2.0%	18.5%
2003	2.8%	29.2%	1.6%	18.3%
2004	3.6%	26.4%	2.1%	16.1%
2005	2.6%	25.1%	1.5%	15.3%

[1] Based on the kernel density estimate of the DSR distribution.

Table 8 shows that the density in the vulnerable tail by income groups. One thing to note is that the results become less robust as the sample size decreases as in the case of the vulnerable density of households by income groups above the 40 per cent IP-DSR threshold. For this reason we focus more on the 23 per cent IP-DSR threshold in Table 8. The right-hand side of the table shows that density of households with an IP-DSR greater than 23 per cent has fallen for both the low and middle income groups since 1999. This is especially apparent for the low-

[33] The 40 per cent threshold is based on gross income and thus does not suffer from this problem.

[34] Based on aggregate data from Statistics Canada's National Balance Sheet Accounts (2005); average for 1999-2005. This is, of course, a simplification of reality as the wedge between gross and disposable income may vary across income groups.

income households, which have the highest density in the vulnerable tail. Meanwhile, the vulnerable tail density for high-income households has remained roughly unchanged.[35]

Table 8

Density in vulnerable tail (indebted households)[1]

	IP-DSR >40%			IP-DSR >23%		
	Low	**Middle**	**High**	**Low**	**Middle**	**High**
1999	5.6%	2.3%	0.7%	37.7%	32.7%	19.0%
2000	8.8%	4.0%	2.0%	39.1%	34.5%	20.3%
2001	7.3%	4.1%	2.2%	37.0%	36.4%	20.2%
2002	7.0%	2.7%	1.1%	33.7%	32.9%	20.8%
2003	5.6%	2.7%	1.3%	39.7%	32.9%	21.0%
2004	7.1%	4.0%	1.5%	32.9%	31.3%	18.9%
2005	5.2%	2.7%	1.4%	32.8%	29.7%	19.4%

[1] Based on the kernel density estimate of the DSR distribution.

Summary

The analysis of the DSR distribution yields the following main results:

- The distribution is asymmetric with a thin, long right-hand tail.
- There has not been any noticeable change in the shape of the distribution since 1999.
- The density of households with a high DSR has fallen since 1999, particularly for the low income group.

Overall, the findings from the distributional analysis of the DSR suggest that Canadian households have fared quite well over the 1999-2005 period even as the sector (as a whole) accumulated debt at a strong pace.

6. Conclusions and future work

An analysis of household indebtedness based solely on aggregate data may be misleading as the aggregate data can mask important information about the distribution of the debt service burden across households. Our examination of the distribution of the DSR across Canadian households for the 1999-2005 period shows that the messages coming from the aggregate and micro data are consistent: household debt burden has fallen over recent years and household financial health remains sound. In particular, debt and asset holdings of households are relatively well matched, the distribution of the IP-DSR is skewed to the left and the shape of the distribution has remained roughly unchanged since 1999.

[35] The qualitative conclusions are the same if we examine the densities for all households, instead of only those with positive debt (as shown in Table 8).

However, it remains important to continue monitoring the distribution of the debt burden, in conjunction with the analysis of the aggregate DSR for households. It is probable that discrepancies between the aggregate DSR and the distribution of the debt burden become more apparent prior to or during periods of asset price misalignment, and other negative macro events. More years of data will be able to provide better insight into this.

Future work on this topic will focus on developing a framework for using the distribution of the DSR for policy analysis simulations. For example, we would like to ascertain how the DSR distribution would behave in response to a monetary policy or an income shock. In addition, we will also construct an aggregate measure of the household debt burden, which includes principal repayment obligations. The aggregate and micro-data based DSR measures will continue to play complementary roles in the Bank of Canada's analysis of household debt.

Appendix 1:
Comparing the IO-DSR from CFM with the aggregate IO-DSR

One check on the quality of the CFM data is to see how well the DSR measure based on these data compares with the measure based on aggregate data. Since the aggregate DSR measure is an interest only measure, we focus on comparing it with the IO-DSR measure from CFM.[36]

The IO-DSR from CFM is calculated using information on the household's loan balances and the corresponding interest rates on the loans:

$$IO-DSR_{t,j} = \frac{\sum_{i} r_{i,j,t} \times loan\_bal_{i,j,t}}{GI_{t,j}} \quad (3)$$

Where:

'*i*' = mortgage loans, personal lines of credit, auto loans, outstanding credit card balance, other personal loans,

'*j*' = household ID,

'*r*' = annualized interest rate on loan as reported by each household,

'*loan_bal*' = outstanding loan balance, and

'*GI*' = gross household income.

The aggregate IO-DSR is estimated as described in *Section 2* of the paper and uses information on the effective interest rate on debt from CFM. The main difference between the aggregate and CFM IO-DSR is, therefore, the debt and income data used in calculating the ratio.

The aggregated IO-DSR from CFM (for all households, including those with zero debt) compares well with the estimate from aggregate data. In particular, both measures show a similar trend of declining debt service burden since 1999 (Figure 4). There is, however, a level difference between the two estimates. Possible reasons for this difference include:

1. the coverage of the debt data: CFM debt levels would be necessarily lower than the aggregate data, given that the aggregate data also include loans to un-incorporated businesses, and
2. the two measures use different denominators in their calculations: the aggregate IO-DSR measure uses disposable income while the CFM based measure used gross. Gross income is (on average) higher than disposable income.

Both these factors would lead to a lower estimate of the IO-DSR from CFM than from the aggregate data. Other factors that could lead to discrepancies between the two estimates include range coding of loan balances and gross income in CFM.[37]

---

[36] CFM allows calculation of both the IO and IP-DSRs.

[37] Loan balances and household income are recorded as ranges and not as a point estimate. For example if household 1 and 2 have gross incomes of $35,001 and $44,001 respectively, they will both fall into the $35,000-$44,999 income group and for our calculations will have an estimated gross income of $40,000. This may lead to some noise in the CFM estimate.

Figure A1

IO-DSR calculations: micro and aggregate data

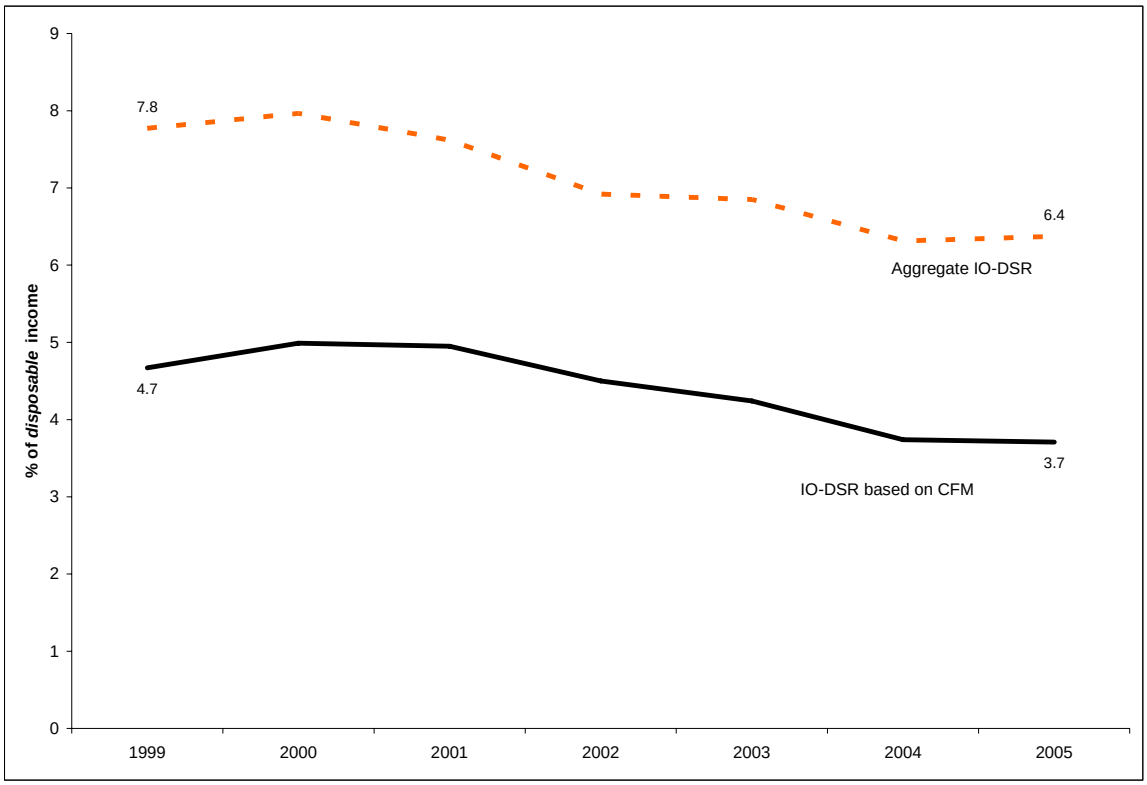

Sources: Statistics Canada, CFM and our calculations.

Appendix 2:
IO- and IP-DSR measures from CFM

CFM allows calculation of both the IO- and IP-DSR for households. Both ratios use gross income in the denominator. The calculation of the IP-DSR for each household is shown by equation 2 in section 5, while the IO-DSR for each household is estimated as follows:

$$IO-DSR_{t,j} = \frac{\sum_i r_{i,j,t} \times loan\_bal_{i,j,t}}{GI_{t,j}} \qquad (4)$$

Where:

'*i*' = mortgage loans, personal lines of credit, auto loans, outstanding credit card balance, other personal loans,

'*j*' = household ID,

'*r*' = annualized interest rate on loan (as reported by the respondent),

'*loan_bal*' = outstanding loan balance, and

'*GI*' = gross household income.

Figure A2 shows that both the IO- and IP-DSRs calculated from CFM[38] suggest a declining debt burden for Canadian households over the 1999-2005 period.

Figure A2

IO- and IP-DSR from CFM (1999-2005)

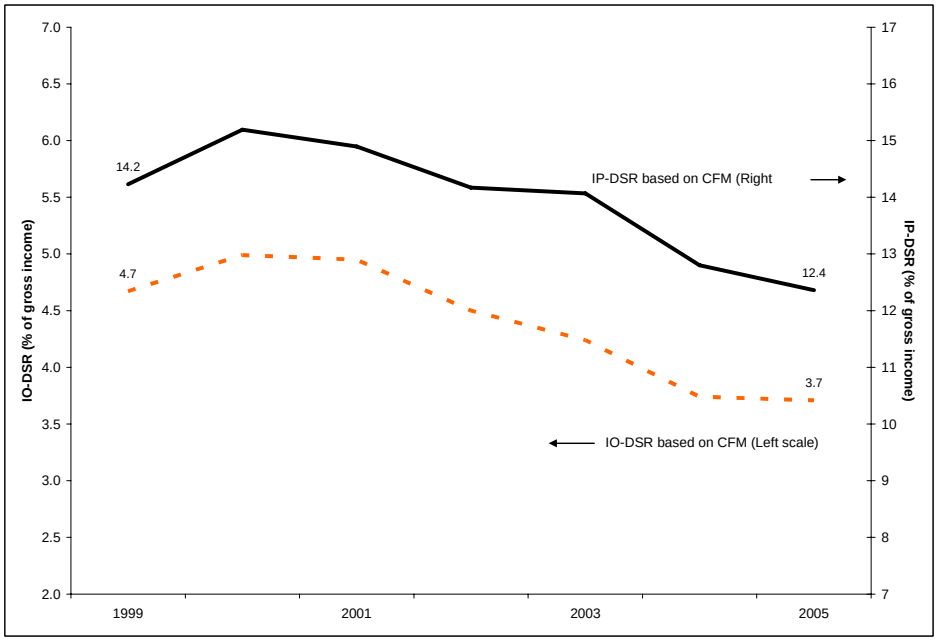

[38] The DSR estimates in figure A2 are estimated as follows. First, a DSR is calculated for each household. The mean DSR is then calculated for each year using CFM weights to aggregate across households. The estimates include all households (for a given year), including those with no debt (i.e. DSR = 0).

Bibliography

Barnes, S. and G. Young (2003): “The rise in US household debt: assessing its causes and sustainability.” Bank of England Working Paper no. 206.

Canner, G.B., A.B. Kennickell and C.A. Luckett (1995): “Household sector borrowing and the burden of debt.” *Federal Reserve Bulletin*. April 1995: 323-338.

DeVaney, S.A. (1994): “Usefulness of financial ratios as predictors of household insolvency.” *Financial Counselling and Planning*, Vol. 5, 1994: 5-24.

Edelberg, W.M. and J.D.M. Fisher (1997): “Household Debt.” *Chicago Fed Letter*. November 199, Number 123.

Herrala, R. (2006): “Household indebtedness.” *Bank of Finland Bulletin*. 2006-1: 91-99.

May, O., M. Tudela and G. Young (2004): “British household indebtedness and financial stress: a household-level picture.” *Bank of England Quarterly Bulletin*. Winter 2004: 414-428.

Reserve Bank of Australia (2003): “Household debt: what the data show.” *Reserve Bank of Australia Bulletin*. March 2003: 1-11.