

## DO THE RICH SAVE MORE? EVIDENCE FROM CANADA

BY SULE ALAN

*University of Essex*

KADIR ATALAY\*

*University of Sydney*

AND

THOMAS F. CROSSLEY

*University of Essex and Institute for Fiscal Studies, London*

This paper provides new evidence on the long-standing question of whether more affluent households save a larger fraction of their income. The major difficulty in empirically assessing the relationship between incomes and saving rates is to construct a credible proxy for long-run income—purged of transitory fluctuations and measurement error. The Canadian Family Expenditure Survey provides us with both unusually good data on savings rates and potential predictors with which we can construct reliable long-run income proxies. Our empirical analysis suggests that the estimated relationship between saving rates and long-run incomes is sensitive to the predictor used to proxy long-run income. Nevertheless, our preferred estimates indicate that, except for poorest households (who simply do not save), saving rates do not differ substantially across predicted long-run income groups.

**JEL Codes:** C81, D12

**Keywords:** income, permanent income, saving rates

### 1. INTRODUCTION

Do the rich save more? This is a critical question for a range of important policy questions. Whether the rich save more often is key determinant of whether a switch from income taxation to consumption taxation is regressive (Crossley *et al.*, 2010). In optimal tax theory, if high ability individuals save more, a non-zero rate of capital taxation can be desirable (Diamond and Spinnewijn, 2011). Whether the rich save more is also relevant to the origins of wealth inequality; the intergenerational transmission of inequality; the consequences of inequality for

*Note:* For helpful comments and suggestions we thank Kevin Milligan, Yuri Ostrovsky, and Garry Barrett, as well as the editor Robert Hill and two anonymous referees. Financial support from the Faculty of Arts and Social Sciences at the University of Sydney, Social and Economic Dimensions of an Aging Population (SEDAP) Research Program at McMaster University, the Social Sciences and Humanities Research Council of Canada (SSHRC), and the ESRC through the ESRC-funded Centre for Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (CPP, reference RES-544-28-5001) is gratefully acknowledged. All errors are our own.

\*Correspondence to: Kadir Atalay, School of Economics, University of Sydney, 2006, NSW, Australia (kadir.atalay@sydney.edu.au).

economic growth; the effect of a tax cut on aggregate demand; and the incidence of the tax expenditures associated with tax-favored saving accounts.

Most non-economists would find the proposition that the rich save more to be obvious. Many economists are more skeptical, for both theoretical and empirical reasons. And despite research spanning more than half a century, the issue is still debated.

The purpose of this paper is to provide new evidence on this long-standing question. This evidence is based on a set of household budget surveys conducted by Statistics Canada in different years. These Family Expenditure Surveys are of unusually high quality and have several unique design features which make them much better suited to the measurement of household saving than similar surveys in other countries. This makes our evidence a valuable contribution to the international debate.

Reasons that many economists doubt that the rich save more include the following. First, since Friedman (1957), economists have emphasized that if agents are forward looking and try to smooth transitory income fluctuations, then a strong correlation between measured current income and saving rates is to be expected. A positive transitory fluctuation raises both current income and current saving. Measurement error in income works in the same way: a positive error in measured income increases both measured saving (measured income minus consumption) and measured income. The implication is that correlation between measured current income and saving rates tells us little about the relationship between true saving rates and average or long-run incomes.

The second reason is the logic of budget constraints: for a given bequest level, a lifetime budget constraint implies that if a lifetime income group saves more rapidly at some ages, the same group must dis-save more rapidly at other ages.

The third reason is that most of our standard models assume features (for example, inter-temporally additive, constant relative risk aversion preferences) that are analytically convenient exactly because of the homotheticity they deliver. So many of our theoretical models “scale” (so that a rich household is identical to several poor households) that the idea that the world also scales has become part of our intuition. Adding idiosyncratic uncertainty to standard saving models can deliver differences in saving rates across long-run income groups but usually in the opposite direction to the non-economists’ intuition: the poor save more. Of course, it is theoretically possible to generate saving rates that increase with lifetime income. Introducing wealth into the utility function in an appropriate way will do this, though such mechanisms are sometimes regarded as artificial. A bequest motive can deliver increasing saving rates with long-run income, if bequests are a luxury good, or if the lifetime earning capacities of successive generations in a dynasty are mean reverting (so that a rich household will expect their children to have lower long-run earnings and a poor household will have the opposite expectation). However, direct empirical evidence on bequests is inconclusive. For example, Hurd (1987) suggests that desired bequests are small in the U.S. while Altonji and Villanueva (2007) conclude the opposite. As both authors acknowledge, research in this area is limited by the available data.

Finally, the proposition that the rich save more is hard to reconcile with long-run trends in saving. Simon Kuznets won a Nobel prize in part for

demonstrating stability of consumption to income ratios over many decades. The average saving rate did not appear to rise as people got richer on average.

Still, aggregate saving rates are affected by many factors. The cross-sectional correlation between current income and savings rates is very strong. And it is true that in a model with idiosyncratic income uncertainty and asset-tested social insurance programs, asset testing can distort saving incentives and lead poor households to save less (as in Hubbard *et al.*, 1994; Huggett and Ventura, 2000). Thus the relationship between affluence, or long-run income and saving rates, remains an important empirical question.

This question was in fact the subject of substantial, if inconclusive, empirical work in the years after Friedman's seminal contribution. For example, Friedman himself found evidence for the "proportionality hypothesis." In contrast, Mayer (1972) found an elasticity of consumption with respect to lifetime income that was less than one. After a period in which the empirical literature on consumption and saving pursued other issues, attention returned to the issue of lifetime incomes and saving rates. Bernheim and Scholz (1993) and Hubbard *et al.* (1994) argue that wealth levels are disproportionately high among households with high long-run or lifetime incomes. On the other hand, Venti and Wise (1998, 1999) and Gustman and Steinmeier (1999) both examine lifetime earnings and retirement wealth in the Health and Retirement Survey, and conclude that wealth-income ratios are fairly constant across lifetime income levels. Wealth levels, of course, reflect not only past rates of active saving but also social security rules and past rates of return.

Dynan *et al.* (2004) use three different U.S. datasets and several different identification strategies to estimate the relationship between saving rates and "permanent" income. They conclude that the evidence supports a positive relationship between saving rates and permanent (or long-run) income. Their evidence has been influential. For example, leading public economics textbooks cite this evidence in discussing the relative merits of income and consumption as a tax base (e.g., Gruber, 2010).

This paper provides new evidence on saving by long-run income groups using Canadian household expenditure data and methods similar to those employed by Dynan *et al.* (2004). By long-run income, we mean income purged of measurement error and transitory fluctuations.<sup>1</sup> A key motivation for this analysis is the unique nature of the Canadian data. Cross-sectional patterns of household saving can be studied using household budget surveys (to construct measures of income minus consumption, or active saving) or with panel data on household wealth (calculating saving as the change in wealth, possibly with a correction for capital gains to give a measure of active saving). Canadian survey data on household wealth do not have a panel component. However, comparisons of national budget surveys suggest that the Canadian Family Expenditure Survey (FAMEX) is of very high quality. Moreover, this survey has several specific features (to be described below) that have lead researchers to believe that it can be the basis for a good measure of active saving, and in particular, is superior to comparable U.S. datasets for this

<sup>1</sup>Dynan *et al.* use "permanent" income to refer to the same concept but as that term sometimes has a specific meaning (in the context of a particular decomposition of an income process), we prefer long-run as a more generic term.

purpose. Partly for these reasons, FAMEX has formed the basis of a number of studies of saving behavior by both Canadian and international authors.<sup>2</sup>

Davies and Burbidge (1994) report a strong correlation between saving rates and current income in these data. However, to the best of our knowledge, this is the first analysis of the relationship between saving rates and long-run income to employ these data.

Our empirical analysis suggests that the estimated relationship between saving rates and long-run income is sensitive to the predictor (or instrument) used to proxy long-run income. Nevertheless, our preferred estimates indicate that, except for poorest quintile of households (who simply do not save), saving rates do not differ substantially across long-run income groups.

The next section describes our data in greater detail. Section 3 outlines our empirical methodology. Our results are presented in Section 4, and Section 5 concludes.

## 2. DATA AND SAMPLE

FAMEX is a full household expenditure survey (collecting information on all categories of expenditure by Statistic Canada) until 1997.<sup>3</sup> Unlike most national expenditure surveys, FAMEX did not have a diary component. Instead, face-to-face interviews were conducted in the first quarter of the year to collect income and expenditure information for the previous year. Thus the 1996 data were collected in January, February, and March of 1997 but refer to the 1996 calendar year. Respondents were asked to consult bills and receipts and if necessary, multiple visits were made to a household. FAMEX was therefore an unusual kind of recall survey in which a considerable effort is made to ensure the quality of the data. Barrett *et al.* (2013) show that aggregates derived from FAMEX had high and stable correspondence to figures from the national accounts, and that response rates to FAMEX did not decline in the 1990s, a period when response rates to household budget surveys in the U.S., U.K., and Australia were declining, in some cases steeply.<sup>4</sup> Moreover, in contrast to the U.S. Consumer Expenditure Survey, FAMEX was particularly designed to capture good quality income information that refers to the same (annual) period as the expenditure information. This facilitates the calculation of household saving.

Our main analysis is based on public use files from the 1996 survey (we also examine earlier surveys to check the robustness of our results). The 1996 survey was chosen because it is the last year in which the principal and interest components of mortgage payments are reported separately. We treat the former as saving and the latter as expenditure.

In studying the relationship between saving rates and long-run incomes, the appropriate saving concept is active saving: the amount of current net income not spent. An important feature of FAMEX is that it contains two measures of

<sup>2</sup>See, for example, Burbidge and Davies (1994), Carroll *et al.* (1994), Davies and Burbidge (1994), Engelhardt (1996), Burbidge *et al.* (1998), and Milligan (2002).

<sup>3</sup>From 1997 Statistics Canada's household budget survey is the Survey of Household Spending.

<sup>4</sup>The response rate to the Canadian household expenditure survey did decline after the FAMEX was replaced by the Survey of Household Spending in 1997.

household active saving. The first is simply after-tax current income (excluding capital gains) minus total expenditure. (This measure is also used in Dynan *et al.* (2004) for the Consumer Expenditure Survey (CEX) and Panel Study of Income Dynamics (PSID); in the case of the latter, total expenditure must be imputed from expenditure on a subset of goods.)

The second measure of active saving in FAMEX is net changes in assets and debts excluding capital gains. This measure is unique (not available in the U.S. surveys). It includes changes in accounts at banks and trust and loan companies; changes in money owed; money deposited as a pledge against future purchases of goods and services; net contributions to and withdrawals from Registered Retirement Saving Plans (a kind of tax-favored individual retirement account); net purchases less sales of financial assets; and sales of personal property.

In the conduct of FAMEX these two measures were partially reconciled in that households in which the two measures show an excessive disparity are asked to review their reports of incomes and expenditures. This “balance edit” has been shown to improve income reports, particularly among low income households (Brzozowski and Crossley, 2011), and thus should improve the saving measure based on income minus expenditure. At the same time, this reconciliation means that measurement errors are unlikely to be independent across the two measures. Following Dynan *et al.*, we divide our saving measures by current income to derive saving rates. The correlation between the two saving rate measures in our data is 0.77.

In summary, there are reasons to believe that the “income minus consumption” measure in FAMEX is superior to those in CEX (where the income data is not ideally suited to this purpose) and PSID (in which total expenditure must be imputed), and FAMEX contains a second measure of annual active saving that can be exploited in a number of ways (which we outline below). Against this, unlike Dynan *et al.*, we do not have panel data on wealth or income. The former is perhaps less important—the annual change in assets and debts recorded in FAMEX provides a second measure of saving that is likely as good as the change in wealth in the Survey of Consumer Finances (SCF) or PSID.<sup>5</sup> However, the absence of panel data on income means that we cannot employ some of the strategies that Dynan *et al.* employ for measuring long-run income.

There are 10,085 respondent households in the 1996 FAMEX. Our estimation sample is restricted in a number of ways. For comparability, we follow the sample selection rules of Dynan *et al.* as closely as possible. First, we exclude households that reported less than \$1000 of income and households that did not report their education level. Second, we delete multiple family units (more than one family living in the same dwelling), which are a small fraction of the sample.<sup>6</sup> Finally, in our analysis we restrict attention to households whose head is between 30 and 59 years of age. The reason for this is to abstract as much as possible from the issues

<sup>5</sup>Although the Survey of Consumer Finances is a very detailed wealth survey, the sample size for the panel component Dynan *et al.* use is very small and subject to a serious attrition problem. The structure of the attrition is documented by Kennickell and Woodburn (1997). The limitations of the PSID wealth data are well known. Changes in the stock of wealth can only be calculated at longer intervals (PSID wealth supplements were conducted at five-year intervals in 1984, 1989, 1994, 1999, and thereafter every two years). In both cases, changes in wealth must be purged of capital gains to construct a measure of active savings.

<sup>6</sup>3 percent of the full sample and 1 percent of the couples only sample.

regarding educational choice and dis-saving in retirement. The resulting sample contains 6062 households. For some of the analysis below we focus on the subset of these households that contain couples (with and without children), which is 4205 households.

As described above, active saving ( $S$ ) is defined as either after-tax net income minus total consumption ( $Y - C$ ) or as the net changes in assets and debts excluding capital gains ( $\Delta A$ ), and then we divide by current income to give the saving rate  $\left(\frac{S}{Y}\right)$ . Income is net household income after taxes and includes wages and salaries, investment income, self-employment earnings, government transfers (Canada or Quebec pension plan benefits, employment insurance benefits, child tax benefits, workers' compensation benefits, goods and services tax credit, provincial tax credits, veterans pension and allowances), and income from other sources (alimony, child support, and income received from tax sheltered saving plans; specifically Registered Retirement Savings Plan (RRSP) annuities received and Registered Retirement Income Fund (RRIF) withdrawals). Total consumption is constructed based on total expenditure and includes expenditures for housing, food, clothing, household operations, personal care, transportation, recreation, education, tobacco and alcoholic beverages, reading materials, and miscellaneous expenses. We treat gifts and contributions to persons outside the household and the interest portion of mortgage payments as consumption. The portion of mortgage payments that is the principal repayment is treated as saving. Note that individual contributions to tax sheltered savings plans (RRSPs) and employer sponsored pensions (Registered Pension Plans, RPPs) are counted as saving, while contributions to public pensions are not (neither employee nor employer contributions to these plans are counted in net income).<sup>7,8</sup>

Descriptive statistics for income and total consumption for our full and couples samples are reported in Table 1. It is clear that the sample of couples has, on average, higher incomes and higher total consumption than the full sample. This reflects the fact that many of the poorest households in our full sample are singles, or households headed by a single adult. When comparing results across the two samples, it will be important to remember that the couples are on average richer than the full sample.<sup>9</sup>

Online Appendix Table A.1 presents descriptive statistics for our saving measures and the composition of saving across the five income quintiles for the full

<sup>7</sup>For a brief overview of the Canadian Retirement System, see OECD (2011).

<sup>8</sup>The appropriate classification of contributions to public pensions depends on whether the plans are fully funded (in which case contributions are clearly savings) or "pay as you go" (in which case contributions are better considered as taxes). The Canada or Quebec pension plan benefits (CPP and QPP) are hybrid plans and hence the appropriate treatment of contributions is unclear. For consistency with previous Canadian research, such as Davies and Burbidge (1994), we have not treated them as savings. However, we have confirmed empirically that our results are robust to including them in savings. Full details are available on request.

<sup>9</sup>Following Dynan *et al.* (2004) in our summary statistics and baseline estimates we do not make any adjustment for household size or composition to income, consumption, or savings (for example, converting to per capita amounts or dividing by an equivalence scale). However, in the online Appendix we examine the robustness of our estimation results to the controlling for additional variables, including household size and composition.

TABLE 1  
DESCRIPTIVE STATISTICS: INCOME AND TOTAL CONSUMPTION, 1996 FAMEX

Variable	Median	Mean	S.D.	Minimum	Maximum
Full sample, 6062 households					
Gross income	51,000	56,482	35,295	500	292,400
Total consumption	34,167	36,978	19,189	-2,139 <sup>a</sup>	196,428
Net income	39,805	43,272	24,450	1,000	248,685
Couples, 4205 households					
Gross income	61,000	66,187	35,024	2,000	292,400
Total consumption	39,134	42,280	18,714	4,717	196,428
Net income	46,627	50,296	23,913	1,000	248,685

*Note:* <sup>a</sup>The data contain a single observation with negative total consumption. This arises because the household sold a vehicle. Reported results include this household in all calculations, but all of our results are robust to the exclusion of this household from the sample.

sample. The first thing to note is that the  $(Y - C)$  measure of saving suggests higher levels and rates of saving in both samples. This is consistent with under-reporting of consumption, or under-reporting of contributions to financial assets and/or retirement of debt, or both. The second thing that we note is that the composition of savings across income levels is similar in both samples. In particular we observe that retirement saving (RRSP contributions, RPP contributions) and mortgage principal payments are the most important components of saving for high current income households. In the lower current income quintiles retirement saving is small and mortgage principal payments negative (indicating increases in mortgage debt). Of course, in these summary statistics income may be confounded by age. We now turn to the methodology of our main analysis which both conditions on age and tries to measure differences by long-run income group.

### 3. EMPIRICAL METHODOLOGY

We wish to estimate the relationship between saving rates and long-run income:

$$(1) \quad \frac{S}{Y} = f(Y^*) + X\beta + e,$$

where  $Y^*$  is long-run income,  $X$  is a set of other determinants of saving behavior (including age), and  $e$  is a disturbance that captures both unmeasured determinants of saving and measurement error in the saving rate. We might formally define “long-run income” as an annuitization of the present value of perceived present and future consumption possibilities. Less formally, we mean household income purged of measurement error and short-run fluctuations as a measure of affluence.<sup>10</sup> To allow for non-linearities in the relationship between

<sup>10</sup>We are also following Dynan *et al.* (2004) and much of the literature on saving (and on inequality) in focusing on realized income, as opposed to full income (and hence taking labor supply as given). This is not entirely satisfactory, as it ignores important variation in resources. In robustness checks reported in the online Appendix, we do control for the number of employed persons in the household; the results are similar to our main findings.

long-run income and saving rates, we parameterize  $f()$  by a set of five dummies capturing the age-conditional quintile of long-run income to which each household belongs ( $X$  does not include a constant).

The key empirical problem we face is that we do not observe true long-run income ( $Y^*$ ). Current income, ( $Y$ ), is a poor proxy for long-run income because either measurement error or the smoothing of transitory income fluctuations will generate a positive relationship between saving and current income and even when there is no relationship with long-run incomes. Our solution, which follows Dynan *et al.* (2004), is a two-stage estimation procedure. In the first stage we construct long-run income proxies by regressing current income on predictors ( $Z$ ) and age group dummies:

$$(2) \quad Y = Z\alpha + X\gamma + u.$$

Predicted values from this regression are then used as our proxy for true long-run income.

$$(3) \quad \hat{Y} = Z\hat{\alpha} + X\hat{\gamma}.$$

We then assign households to age-conditional predicted long-run income quintiles, and construct the quintile dummies that were described above. In the second stage we estimate equation (1) by quantile regression.<sup>11</sup> Since long-run income is estimated in the first stage we bootstrap the standard errors.

The key to our empirical strategy is obviously the predictors for true long-run income. These must be (i) strongly correlated with true long-run income, but not with the transitory components of current income, and (ii) excludable from the saving equation (uncorrelated with unmeasured determinants of saving and with measurement error in the saving rate). We consider two predictors for true long-run income that are also employed in the literature and by Dynan *et al.*: education and non-durable consumption (or components of non-durable consumption). As noted above, our cross-sectional data does not allow us to use follow Dynan *et al.* in using lags and leads of income as alternative predictors.<sup>12</sup> Both education and non-durable consumption are strongly correlated with true long-run income. However, the second condition—that they are excludable from the saving equation—may be violated for reasons specific to each predictor.

Although it is highly correlated with long-run income, education may also be correlated with unobserved taste variables that, in turn, influence saving behavior. For example, it is plausible to think that educational choices are associated with individuals' planning horizons or risk tolerance; patience and risk tolerance are also associated with higher long-run saving. Note that education could be related to unobserved taste variables because preferences determine education choices or because education affects preferences (such as the planning horizon/discount rate

<sup>11</sup>We also employed mean regression as a robustness check and find that our results are robust to the choice of estimator. These additional results are available on request.

<sup>12</sup>These are not necessarily superior instruments. For example, depending on the income process, they may not be excludable from the saving equation. It would nevertheless be desirable to compare results based on the full range of predication strategies.

or risk tolerance).<sup>13</sup> If education is related to preference heterogeneity that is important for saving behavior, then it is not a valid predictor (because it is not excludable from equation (1)). The likely consequence of this is an upward bias in the estimated relationship between long-run income and saving rates: those with more education are more patient and save more.

The information on education in FAMEX is categorical (less than nine years' education; some or completed secondary education; less than post-secondary; post-secondary education; college degree or higher) and is available for heads and spouses. To maximize the variation in long-run household income that we capture with education, we construct a set of dummies capturing different combinations of head and spouse education observed in households. Consequently, our results using education as a predictor are for the sample of couples only.

We define non-durable consumption as total consumption minus spending on shelter, vehicles, and household furnishings.<sup>14</sup> If we use non-durable consumption ( $C^N$ ) to proxy true long-run income, and total consumption in the calculation of saving rates, then any measurement error that is common to both will enter on both the left side and right side of our estimating equation, and bias our estimates.

$$(4) \quad \frac{Y - C}{Y} = f(C^N \hat{\alpha} + X \hat{\gamma}) + X\beta + e.$$

Long-run income is positively correlated with non-durable consumption, and consumption enters the saving rate negatively. If the true relationship between saving rates and long-run income is positive, then measurement error common to non-durable and total consumption will impart a negative bias to our estimates, biasing them toward zero. The same problem arises if consumption has a transitory component (if some households are liquidity constrained, or because of purchase infrequency).

Fortunately, the data afford us ways of addressing this problem. First, and uniquely with FAMEX, we can replace  $(Y - C)$  by our second measure of saving,  $(\Delta A)$  in our estimating equation. To the extent that measurement errors in  $(\Delta A)$  are not perfectly correlated with measurement errors in consumption, this should reduce the bias.<sup>15</sup> Second, rather than use (all) non-durable consumption as our predictor, we can use components of non-durable consumption that are well-measured and sensitive to long-run income. The second criterion points toward luxury goods. Food expenditure is thought to be well-measured in surveys but food is a necessity. Expenditures that are

<sup>13</sup>Bernheim *et al.* (2001) provide some evidence that consumer studies in high schools raise future saving rates.

<sup>14</sup>Thus it contains spending on food, household operations, clothing, health care, personal care, tobacco and alcoholic beverages, reading materials, and miscellaneous expenses, plus transportation and recreation minus purchases of cars and recreational vehicles.

<sup>15</sup>An alternative way to exploit the two measures of active saving available for each household is to pool the data and treat it as a panel with two observations on each household (these are repeated measures, but not temporally separated—they refer to the same year). In principle this could increase the precision of our estimates. We have implemented this strategy, allowing for a common mean shift between the two measures, and taking care in our bootstrapping to resample households (pairs of observations) in order to preserve the correlation structure in the data (as in panel data bootstrapping). In practice, this did not lead to much increase in precision. The results are reported in a working paper version of this paper (Alan *et al.*, 2006).

TABLE 2  
EXPLORATION OF INCOME PREDICTORS

First stage regression:

$$Y = X\alpha + W\beta + Z\gamma + u$$

Income ( $Y$ ) on age dummies ( $X$ ), (additional controls ( $W$ )), and instruments ( $Z$ ), (couples sample)

Partial  $R^2$

	Base +	Base, Food +	Base, Controls +	+ Controls + Food
	$\frac{R^2_{X,Z} - R^2_X}{1 - R^2_X}$	$\frac{R^2_{X,Z_1,Z_2} - R^2_{X,Z_1}}{1 - R^2_{X,Z_1}}$	$\frac{R^2_{X,W,Z} - R^2_{X,W}}{1 - R^2_{X,W}}$	$\frac{R^2_{X,W,Z_1,Z_2} - R^2_{X,W,Z_1}}{1 - R^2_{X,W,Z_1}}$
Total consumption	0.5019		0.4536	
Non-durable cons.	0.4817		0.4328	
Food	0.2467		0.2056	
Phone	0.0595	0.0240	0.0507	0.0225
Water, fuel, and elec.	0.0859	0.0433	0.0667	0.0463
Gas	0.0662	0.0273	0.0359	0.0177
Clothing	0.2954	0.1535	0.2429	0.1341
Education dummies	0.1350		0.1254	

Notes: The  $R^2$  for the age dummies alone is 0.0101. The additional controls are household size, number of children, and number of employed persons in the household. The  $R^2$  for the age dummies plus additional controls is 0.1422.

regularly billed (so that respondents are able to consult bills during the survey interview) are also attractive, so long as expenditure is discretionary.

Table 2 evaluates a number of components of non-durable consumption as income predictors. We assess the marginal contribution of each predictor the first stage regression fit (equation (2)) with the partial  $R^2$ . There are two specifications in the table. The first one is our base specification which includes only age dummies. The second specification adds controls for household size, number of children, and number of employed persons in the household. For comparison purposes we report the partial  $R^2$  belongs for total consumption and for non-durable consumptions. Both explain about half of the remaining variation in income after controlling for age (and other household characteristics, in specification 2).<sup>16</sup> We then turn to potentially better measured subcomponents of non-durable expenditure, including food, phone bills, utility bills, petrol, and clothing. With respect to phone bills, we note that the long distance component of the phone bills is discretionary and potentially a luxury good.

The results suggest that food and clothing both have strong marginal explanatory power, explaining about half of the variation in income that remains after controlling for age or age and other household characteristics. Both food and clothing have greater marginal explanatory power for income than our full set of education dummies; the latter explain less than 15 percent of the residual variation in income. Clothing has a higher partial  $R^2$  than food, despite a lower budget

<sup>16</sup>A possible concern with non-durable consumption as an income predictor is that durables are often luxuries and so non-durable consumption will not capture the full range of income variation. There is not much evidence for that here: non-durable consumption explains almost as much of the variation in income as total consumption.

share, because it is a more luxurious category of expenditure (estimated income elasticities of clothing expenditure are typically above one). How well clothing expenditures are measured, relative to food, is not known. Telephone and utility bills are likely to be well measured (as respondents can consult their bills). The results in Table 2 suggest that such expenditures have less explanatory power than food and clothing, but do have useful explanatory power, either on their own or in conjunction with food expenditure. Based on these results we chose to employ, as predictors for income, three different subcomponents of non-durable consumption: food, food plus clothing, and food plus telephone bills. Again the hope is that these components are well measured, and hence less correlated than non-durable consumption with any measurement error in savings.

Of course, it is unlikely that we can eliminate all bias. What we can do, however, is to assess how serious the bias may be by observing how the estimated relationship between saving rates and long-run incomes changes as we make these substitutions. If measurement error in total and non-durable consumption imparts a significant negative bias to our estimated relationship between saving rates and long-run incomes, then we would expect the estimated relationship to become steeper as we replace  $(Y - C)$  by  $(\Delta A)$ , or replace non-durable consumption by well-measured components of non-durable consumption. We lean heavily on this idea in assessing our results, which are presented in the next section.

## 4. RESULTS

### 4.1. *Savings Rates and Measured Current Income*

Recall that in all our median regressions we suppress the constant and include dummies for all five current (or predicted long-run) income quintiles. Among the (household head's) age dummies we exclude the 40–49 year-old group. Thus, the estimated coefficient on a given income quintile dummy corresponds to the median saving rate of households in that current (or predicted long-run) income quintile whose head is between 40 and 49 years old.

We begin our analysis by documenting the estimated relationship between saving rates and current incomes. Online Appendix Table A.2 presents the results for both the full sample and the sample of couples (with and without children). Results are presented for both saving rate measures and also for RRSP contribution rates and mortgage principal payment (MPP) rates. These results are also summarized in Figure 1. Here, for each set of results, we plot the estimated median saving rate for each current income quintile, against the median income within the quintile. Thus there are two panels (full sample and couples), each with two lines (corresponding to estimates based on  $(Y - C)/Y$  and  $(\Delta A)/Y$ ) and five points on each line (corresponding to the five income quintiles).

The results confirm that saving rates are strongly increasing with current income.<sup>17</sup> For example, focusing on the full sample and the  $(\Delta A)/Y$  saving

<sup>17</sup>We also estimated *marginal* propensity to save (MPS) using the variation in median saving rates and incomes across income quintiles. The MPS ranges from 6 cents per dollar of income between quintiles 1 and 2 up to 30 cents between quintiles 4 and 5. The full results are available on request from authors.

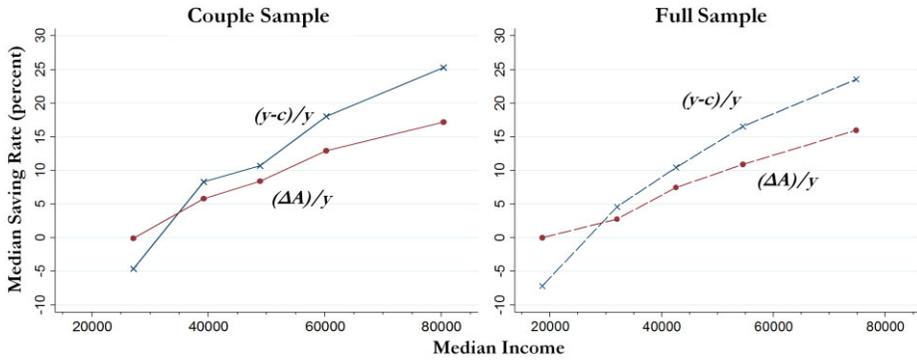


Figure 1. Median Saving Rates and Current Income

measure, median saving rates for 40–49 year-old households range from 0 percent in the lowest income quintile to 16 percent in the highest quintile; the corresponding numbers are –7.2 percent to 23.5 percent when the  $(Y - C)/Y$  measure of saving is used. Using similar methods, Dynan *et al.* report a wider range of estimated savings rates by current income quintiles in the U.S. CEX (–23 to 46 percent); of course, current incomes are more dispersed in the U.S. data. The stars on quintile 2 through 5 coefficients in online Appendix Table A.2 indicate that each coefficient is statistically significantly different from the coefficient for the quintile below it (at the 5 percent level). As in the unconditional numbers reported in Table 2, the RRSP and MPP contribution rates are only positive for the top two income quintiles, and for the lower quintiles median contribution rates are 0 percent.

#### 4.2. Savings Rates and Long-Run Income

We now turn to the relationship between saving rates and long-run incomes, which is our primary interest. Table 3 reports estimated saving rates by predicted long-run income quintiles (from median regressions.) All the estimates in this table are based on the couples sample. The first two columns of Table 3 (on the left) give results using the education of the head and spouse as predictors for long-run income. The two columns report estimates based on  $(Y - C)/Y$  and  $(\Delta A)/Y$  (moving from left to right). These results are summarized in panel (b) of Figure 2. The format of Figure 2 (and subsequent figures) is the same as that of Figure 1 except that each point represents a quintile of predicted long-run income (rather than reported current income). The last two columns of Table 3 (on the right) give results using non-durable consumption as the predictor (with saving measures based on  $(Y - C)/Y$  and  $(\Delta A)/Y$ ). These results are summarized in panel (a) of Figure 2. Panel (c) of Figure 2 also summarizes results using non-durable consumption as the predictor, but for the full sample. (The estimates underlying the  $(\Delta A)/Y$  line in this graph are given in the sixth column of Table 4.)

The first aspect of these results to note is that the choice of saving measure ( $(Y - C)/Y$  or  $(\Delta A)/Y$ ) makes little difference to our central question. Estimates

TABLE 3

MEDIAN REGRESSION OF SAVING RATES ON AGE AND PREDICTED LONG-RUN INCOME QUINTILE DUMMIES (PREDICTING LONG-RUN INCOME WITH EDUCATION OR NON-DURABLE CONSUMPTION; COUPLES SAMPLE)

Predictor(s) for Long-Run Income	Education of Head and Spouse		Non-Durable Consumption	
	(1) $(Y - C)/Y$	(2) $(\Delta A)/Y$	(3) $(Y - C)/Y$	(4) $(\Delta A)/Y$
Quintile 1	8.03 (0.95)	5.72 (0.93)	10.93 (1.29)	6.48 (0.94)
Quintile 2	9.65* (1.33)	6.91 (0.95)	14.88* (1.27)	10.42* (0.84)
Quintile 3	12.35* (1.15)	9.03* (1.01)	11.26* (1.31)	9.17 (0.90)
Quintile 4	16.12* (1.21)	11.98* (0.90)	12.77 (1.18)	9.82 (0.88)
Quintile 5	18.71* (1.42)	14.28 (0.91)	7.88* (1.38)	7.47 (0.93)
Age 30–39	-0.90 (1.01)	0.07 (0.88)	-0.17 (1.10)	0.03 (0.83)
Age 50–59	2.70 (1.15)	0.18 (0.97)	3.01 (1.34)	0.88 (0.91)

Notes: Sample size: 4205 households. Standard errors based on 999 bootstrap replications.

\*The coefficient on this quintile is statistically different than the coefficient on the preceding quintile at the 5% level.

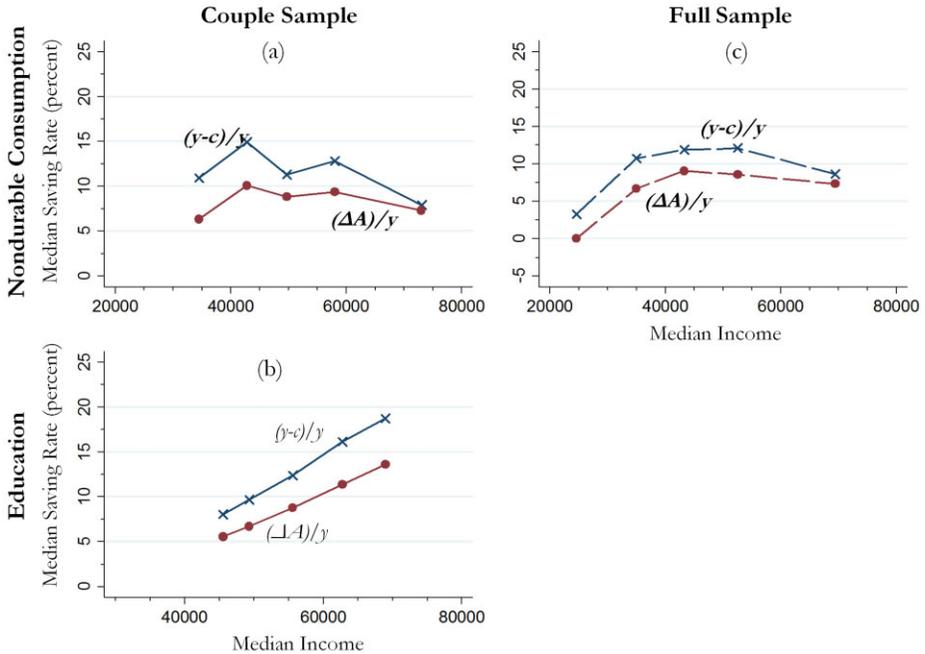


Figure 2. Median Saving Rates and Predicted Long-Run Income Quintiles

based on  $(Y - C)/Y$  give higher saving rates in every quintile than those based on  $(\Delta A)/Y$ . However, the pattern across quintiles is quite similar regardless of choice of measure. In what follows, we focus on the estimates based on  $(\Delta A)/Y$ .

Using education as the predictor for long-run income results in a strong positive relationship between saving rates and predicted long-run incomes. The estimated median saving rate for a 40–49 year-old household rises monotonically from 5.7 percent in the bottom quintile of lifetime incomes to 14.3 percent in the top quintile (2nd column of Table 3 and Figure 2(b)). While only third and fourth quintiles' coefficients are statistically different from one just below them, the coefficient on the top quintile dummy is strongly statistically different from the coefficient on the bottom quintile dummy.

In contrast, when we use non-durable consumption as a predictor, the estimated relationship between saving rates and predicted long-run incomes is essentially flat. The estimated median saving rate for a 40–49 year-old household is 6.5 percent in the bottom quintile of lifetime incomes and 7.5 percent in the top quintile. It actually peaks (at 10.4 percent) in the 2nd quintile.

Figure 2 illustrates an important distinction between the sample of couples (with and without children) and the full sample. For the couple sample, the median predicted long-run income in the bottom quintile is 33,288 CAD (1996 Canadian dollars). For the full sample, the corresponding amount is 22,367 CAD. This is because many of the additional households in the full sample (singles, and single adult headed households) are poorer than those in the couple sample. When we included these poorer households in our estimates, we see a much lower saving rate in the lowest quintile of the predicted long-run income distribution. Using the  $(\Delta A)/Y$  measure, the estimated median saving rate for a 40–49 year-old household in the bottom quintile of predicted long-run incomes in the full sample is 0 (panel (c) of Figure 2 and column 6 of Table 4). Above the first quantile however, the estimated relationship is flat in this sample as well.

The flatness of the relationship between saving rates and predicted long-run incomes when we use non-durable consumption as a predictor for true long-run incomes is consistent with the U.S. evidence based on CEX reported by Dynan *et al.* Although it may be attributed to a downward bias (resulting from measurement error in consumption) in the U.S. study, this seems less plausible here, given the quality of our data and the fact that we obtained the same result when we use the  $(\Delta A)/Y$  measure (which should suffer from less bias.)

To push this further, we replace non-durable consumption as our predictor with components of non-durable consumption. As described in the previous section, this should further reduce potential correlation between measurement errors on the left and right sides of our estimating equation. The results are presented in Table 4 (for both couples and the full sample) and summarized in Figure 3 (for the full sample) and Figure 4 (for the couples). All of these estimates use  $(\Delta A)/Y$  as the measure of the saving rate. In Figure 3 we also include the estimated relationship between saving rates and current incomes for comparison. In Figure 4 we include, for comparison, the estimated relationship between saving rates and current incomes, and the estimated relationship between saving rates and predicted long-run incomes when education is used as the predictor.

TABLE 4  
 PREDICTING LONG-RUN INCOME WITH ALTERNATIVE CONSUMPTION MEASURES (MEDIAN REGRESSION OF SAVING RATES ON AGE AND PREDICTED LONG-RUN INCOME  
 QUINTILE DUMMIES;  $(\Delta A)/Y$  SAVING MEASURE)

Predictor(s) for Lifetime Income	Couples (4205 Households)					Full Sample (6062 Households)				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Non-Durable Consumption	Food	Food and Telephone	Food and Cloth	Total Consumption	Non-Durable Consumption	Food	Food and Telephone	Food and Cloth	Total Consumption
Quintile 1	6.48 (0.94)	7.75 (0.91)	7.62 (0.85)	6.85 (1.00)	8.79 (0.87)	0 (0.15)	0 (0.21)	0 (0.21)	0 (0.12)	0 (0.19)
Quintile 2	10.42* (0.84)	9.68 (1.00)	10.47 (0.98)	9.93* (0.99)	11.7 (0.85)	6.87* (0.75)	5.78* (0.78)	6.39* (0.77)	6.02* (0.91)	8.83* (0.70)
Quintile 3	9.17 (0.90)	9.92 (0.79)	9.87 (0.84)	9.09 (0.78)	10.46 (0.93)	9.35* (0.68)	9.68* (0.76)	9.12* (0.81)	9.75* (0.73)	9.89 (0.64)
Quintile 4	9.82 (0.88)	7.76* (1.05)	8.51 (0.87)	8.52 (0.94)	8.1 (0.82)	8.86 (0.58)	9.05 (0.63)	9.24 (0.59)	8.18 (0.61)	8.61 (0.68)
Quintile 5	7.47 (0.93)	7.25 (0.91)	6.64* (1.00)	8.37 (1.00)	1.18* (1.01)	7.75 (0.63)	7.17 (0.70)	6.89 (0.75)	7.93 (0.71)	2.81* (0.74)
Age 30-39	0.03 (0.83)	0.60 (0.85)	0.51 (0.85)	0.59 (0.81)	0.44 (0.78)	0 (0.27)	0 (0.38)	0 (0.37)	0.06 (0.35)	0.13 (0.39)
Age 50-59	0.88 (0.91)	1.28 (0.93)	1.25 (0.89)	1.45 (0.98)	1.17 (0.99)	0 (0.30)	0.02 (0.47)	0 (0.44)	0 (0.42)	0.16 (0.56)

Notes: Standard errors based on 999 bootstrap replications.  
 \*The coefficient on this quintile is statistically different than the coefficient on the preceding quintile at the 5% level.

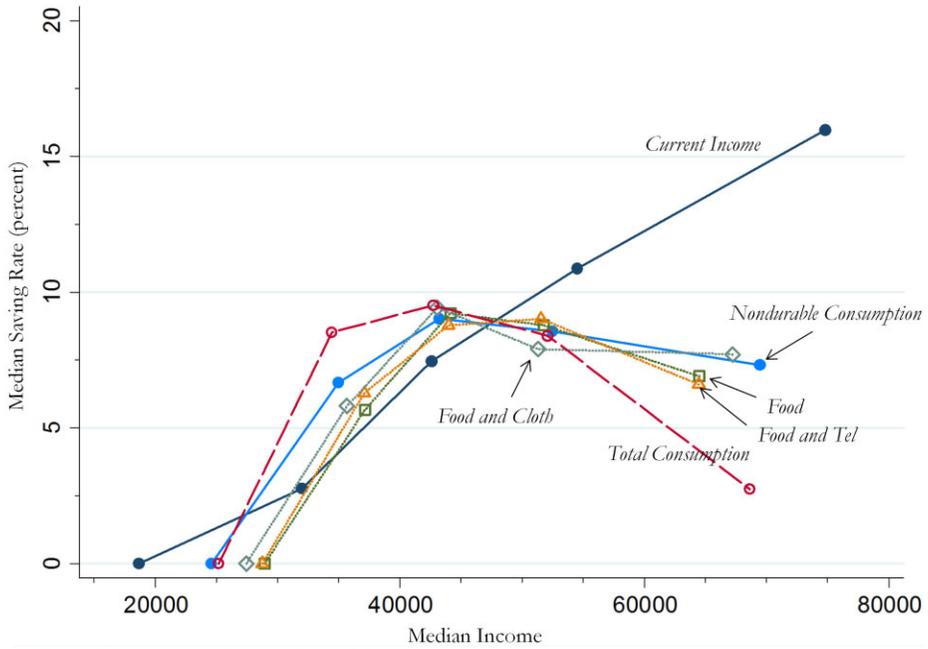


Figure 3. Median Saving Rates and Predicted Long-Run Income Quintiles, Full Sample,  $(\Delta A)/y$  Saving Measure, Alternative Predictors

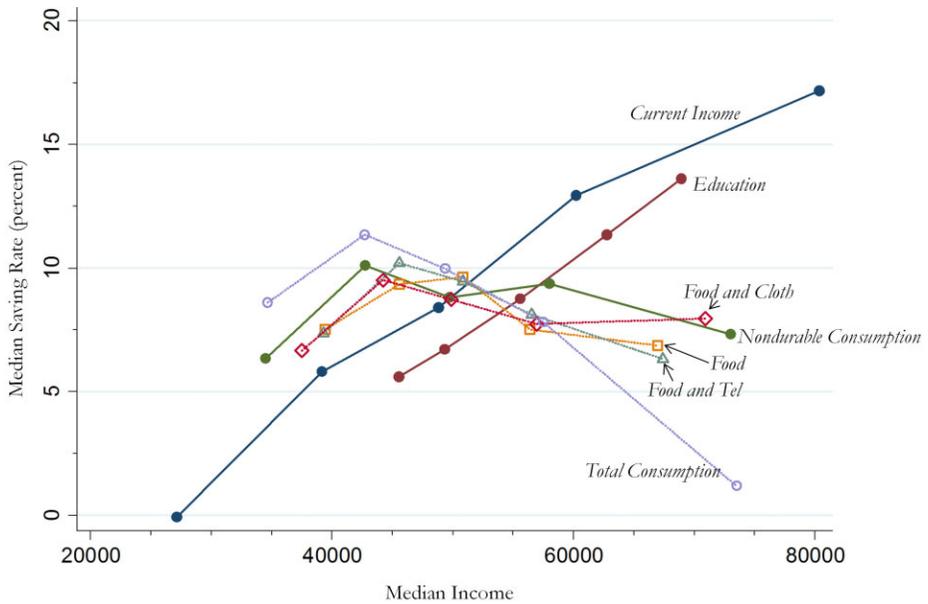


Figure 4. Median Saving Rates and Predicted Long-Run Income Quintiles, Couples Sample,  $(\Delta A)/y$  Saving Measure, Alternative Predictors

The main message of these results is that estimated relationship between saving rates and long-run incomes is not sensitive to whether we use non-durable consumption as a predictor or a component of non-durable consumption. If we use the sample of couples (with and without children) the estimated relationship is essentially flat. If we use the full sample, so that the bottom quintile is poorer, we see low saving in the bottom quintile, and then a flat relationship in the next four quintiles. Using any non-durable consumption measure as a predictor for long-run income results in an estimated relationship between saving rates and long-run incomes that is much flatter than the estimated relationship between saving rates and current incomes. In contrast, when we use education as a predictor for long-run incomes, the estimated relationship between saving rates and long-run incomes is as steep as the estimated relationship between saving rates and current incomes (see Figure 4).

Note also that if we use total consumption as our predictor of long-run income, so that we have the same measure of consumption on the left- and right-hand sides of our estimating equation, we find a strongly negative relationship between savings and long-run incomes (Figures 3 and 4). Our interpretation is that the negative bias arising from measurement error or transitory consumption is important in this case, but not if we use non-durable consumption or well-measured subcomponents of non-durable consumption.

## 5. CONCLUSION

To summarize, when we use education as a predictor for long-run income, we find a strong positive relationship between saving rates and predicted long-run income. Indeed, these results suggest that the relationship between saving rates and predicted long-run income is just as steeply positive as the relationship between saving rates and current income. This would be surprising, as it is likely that at least part of the correlation between saving rates and current incomes reflects the smoothing of transitory income shocks or measurement error. In other words, to give credence to these estimates seems to imply that current measured income is measured without error and that saving rates do not respond to transitory fluctuations in income (households do not smooth).

In contrast, when we use non-durable consumption as a predictor for long-run income, we find that above the bottom predicted long-run income quintile, saving rates are fairly flat. A concern with these results is that measurement error in consumption imparts a negative bias to the estimated relationship. However, when we take steps to mitigate this bias (constructing saving from net changes in assets and debts rather than income minus consumption; and using well-measured components of non-durable consumption as predictors) we observe very little change in the estimated relationship. This stability of the estimates as we take progressive steps to eliminate the putative bias is inconsistent with the view that measurement error in consumption imparts a substantial negative bias to the estimates.

Consequently, we believe that the most reasonable interpretation of the data is that education is a poor predictor in this context, not because it does not predict long-run incomes but because it is correlated with unobserved tastes for saving. The best guide to the relationship between saving rates and long-run incomes is the

estimates that use consumption as a predictor for long-run income. We therefore conclude that the rich do not save more—at least compared to those in the middle of the predicted long-run income distribution. Saving rates are very flat above the bottom quintile of predicted long-run incomes. However, the poor—those in the bottom quintile—save very little.

This conclusion differs somewhat from that reached by Dynan *et al.* (2004), largely because we put much greater weight on the results that use non-durable consumption as a predictor. We are able to do so because of the quality and unique features of the FAMEX data. In a recent paper, Bozio *et al.* (2011) applied a similar methodology to U.K. data and found that saving rates rise with education. They do not report estimates where long-run incomes are proxied by non-durable consumption, citing the concern that measurement error in consumption would downward bias the estimates. But like the U.S. consumption data, the U.K. data lack the features of the Canadian data that make this strategy credible. And while our conclusions are different from Dynan *et al.* (2004), they are in accord with other U.S. research using different data and methods (Venti and Wise, 1998, 1999; Gustman and Steinmeier, 1999).

It is important to note that our analysis, and that of the papers above, which employ survey data, is not informative about saving behavior of the very rich (the top 1 percent). These households are not well captured by survey data. There is extensive evidence, reviewed by Carroll (2000), that these households are substantially different than the rest of the population, and that, in particular, makes it hard to reconcile their saving behavior with standard life-cycle models of saving.

Thus our preferred reading of the evidence is that standard economic models of saving (which by and large imply constant saving rates by lifetime income) might provide reasonable guidance to the types of policy questions raised in the introduction—except when those questions pertain to the poorest quintile and to the very rich.

Our results confirm that those that are poor in a lifetime sense do not save: it is not just the case that those with transitorily low income dis-save. This suggests an important priority for future research. Are the low savings rates of the poor a rational response to disincentives in social insurance programs (as suggested by Hubbard *et al.*, 1994)? Social insurance programs may discourage the saving of poorer households in two ways. First, the insurance provided by these programs may diminish the precautionary saving motive (“crowding out” self-insurance). Second, the means-testing and claw-backs in such programs may mean that the poor face very low after-tax returns on saving. Shillington (2003) has pointed out that the combination means testing of income support and income taxes will result in many Canadian seniors of modest means facing tax rates of 100 percent or more on income tax-favored retirement saving accounts (RRSPs). Thus, these households may have very little retirement saving motive. Alternatively, do the very low savings rates of the poor reflect something about preferences of fundamental constraints (such as a “consumption floor”) or about other aspects of behavior (such as a limited capacity to plan or optimize)? The policy implications of these alternative explanations are very different and hence further research to fully establish the role of each in shaping the saving behavior of lower income households remains important.

## REFERENCES

- Alan, S., K. Atalay, and T. F. Crossley, "Do the Rich Save More in Canada," *Social and Economic Dimensions of an Aging Population Research Papers 153*, McMaster University, 2006.
- Altonji, J. and E. Villanueva, "The Marginal Propensity to Spend on Adult Children," *BE Journal of Advances in Economic Analysis and Policy*, 7, 2007.
- Barrett, G., P. Levell, and K. Milligan, "A Comparison of Micro and Macro Expenditure Measures Across Countries Using Differing Survey Methods," in C. Carroll, T. Crossley, and J. Sabelhaus (eds), *Improving the Measurement of Expenditure Measures*, NBER-CRIW Conference Volume, 2013.
- Bernheim, B. D. and J. K. Scholz, "Private Saving and Public Policy," in J. M. Poterba (ed.), *Tax Policy and the Economy*, MIT Press, Cambridge, MA, Vol. 7, 73–110, 1993.
- Bernheim, B. D., D. M. Garrett, and D. M. Make, "Education and Saving: The Long-Term Effects of High School Financial Curriculum Mandates," *Journal of Public Economics*, 80, 435–65, 2001.
- Bozio, A., C. Emmerson, C. O’Dea, and G. Tetlow, "Do the Rich Really Save More. Evidence from Lifetime Earnings and Consumption Data in the UK," Mimeo, 2011.
- Brzozowski, M. and T. F. Crossley, "Measuring the Well-Being of the Poor with Income or Consumption: A Canadian Perspective," *Canadian Journal of Economics* 44(1), 88–106, 2011.
- Burbidge, J. B. and J. B. Davies, "Government Incentives and Household Saving in Canada," in J. M. Poterba (ed.), *Public Policies and Household Saving*, University of Chicago Press for the NBER, 19–56, 1994.
- Burbidge, J. B., D. Fretz, and M. R. Veall, "Canadian and American Saving Rates and the Role of RRSPs," *Canadian Public Policy*, 24, 259–63, 1998.
- Carroll, C. D., "Why Do the Rich Save So Much?" in J. B. Slemrod (ed.), *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*, Harvard University Press, Cambridge, MA, 2000.
- Carroll, C. D., B.-K. Rhee, and C. Rhee, "Are There Cultural Effects on Saving? Some Cross-Sectional Evidence," *Quarterly Journal of Economics*, 109, 685–99, 1994.
- Crossley, T. F., D. Phillips, and M. Wakefield, "Value Added Tax," in R. Chote, C. Emmerson, and J. Shaw (eds), *The IFS Green Budget 2009*, The Institute for Fiscal Studies, London, 2010.
- Davies, J. B. and J. B. Burbidge, "Household Data on Saving Behavior in Canada," in J. M. Poterba (ed.), *International Comparisons of Household Saving*, University of Chicago Press for the NBER, Chicago, IL, 11–56, 1994.
- Diamond, P. and J. Spinnewijn, "Capital Income Taxes with Heterogeneous Discount Rates," *American Economic Journal: Economic Policy*, 3(4), 52–76, 2011.
- Dynan, K. E., J. Skinner, and S. P. Zeldes, "Do the Rich Save More?" *Journal of Political Economy*, 112, 397–444, 2004.
- Engelhardt, G. V., "Tax Subsidies and Household Saving: Evidence from Canada," *Quarterly Journal of Economics*, 111, 1237–68, 1996.
- Friedman, M., *A Theory of Consumption Function*, Princeton University Press, Princeton, NJ, 1957.
- Gruber, J., *Public Finance and Public Policy*, 3rd edition, Worth Publishers, New York, 2010.
- Gustman, A. and T. Steinmeier, "Effects of Pensions on Savings: Analysis with Data from the Health and Retirement Study," *Carnegie-Rochester Conference Series on Public Policy*, 50, 271–324, 1999.
- Hubbard, R. G., J. Skinner, and S. P. Zeldes, "Precautionary Saving and Social Insurance," *Journal of Political Economy*, 103, 360–99, 1994.
- Huggett, M. and G. Ventura, "Understanding Why High Income Households Save More than Low Income Households," *Journal of Monetary Economics*, 45, 361–97, 2000.
- Hurd, M. D., "Savings of the Elderly and Desired Bequests," *American Economic Review*, 77, 298–312, 1987.
- Hurst, E. and A. Lusardi, "Liquidity Constraints, Wealth Accumulation, and Entrepreneurship," *Journal of Political Economy*, 112, 319–47, 2004.
- Kennickell, A. B. and L. Woodburn, "Weighting Design for the 1983–89 SCF Panel," Federal Reserve Board of Governors, Washington DC, 1997.
- Mayer, T., *Permanent Income, Wealth, and Consumption: A Critique of the Permanent Income Theory, the Life-Cycle Hypothesis, and Related Theories*, University of California Press, Berkeley, CA, 1972.
- Milligan, K., "Tax-Preferred Savings Accounts and Marginal Tax Rates: Evidence on RRSP Participation," *Canadian Journal of Economics*, 35, 436–56, 2002.
- OECD, "Pensions at a Glance 2011: Retirement-Income Systems in OECD and G20 Countries" ([www.oecd.org/els/social/pensions/PAG](http://www.oecd.org/els/social/pensions/PAG)), 2011.
- Shillington, R., "New Poverty Traps: Means-Testing and Modest-Income Seniors," C.D. Howe Institute Backgrounder No. 65, C.D. Howe Institute, 2003.

- Venti, S. and D. Wise, "The Cause of Wealth Dispersion at Retirement: Choice or Chance?" *American Economic Review*, 88, 185–91, 1998.
- , "Lifetime Income, Saving Choices, and Wealth at Retirement," in J. Smith and R. Willis (eds), *Wealth, Work, and Health. Innovations in Measurement in the Social Sciences: Essays in Honor of F. Thomas Juster*, University of Michigan Press, Ann Arbor, MI, 1999.

## SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article at the publisher's web-site:

- Table A.1:** Descriptive Statistics Saving, Saving Components and Saving Rates, 1996
- Table A.2:** Median Regression of Saving Rates on Age and Current Income Quintile Dummies
- Figure A.1:** Robustness to Including Additional Covariates in to the Estimation
- Figure A.2:** Alternative Survey Years
- Figure A.3:** Non-salary Earners
- Figure A.4:** Older Sample (Age 65+)
- Figure A.5:** Education Expenses as Saving