ESSAYS ON HOUSEHOLD SAVINGS AND FINANCES

ESSAYS ON HOUSEHOLD SAVINGS AND FINANCES

By

KADIR ATALAY, B.S., M.A.

A Thesis

Submitted to the School of Graduate Studies

in Partial Fulfillment of the Requirements

for the Degree

Doctor of Philosophy

McMaster University

DOCTOR OF PHILOSOPHY (2009)
(Economics)

McMaster University
Hamilton, Ontario

TITLE: Essays on Household Savings and Finances

AUTHOR: Kadir Atalay, B.S.

Hacettepe University, Ankara, Turkey

M.A.

York University, Toronto, Ontario

SUPERVISOR: Professor Paul Contoyannis

NUMBER OF PAGES:

xii, 198

Abstract

My doctoral dissertation is composed of an introductory chapter followed by five independent chapters on household savings and finances.

After the introductory chapter, the second chapter investigates the living standards of Canadian retirees and the adequacy of their financial preparations. We explore the responses of Canadian retirees to subjective survey questions administrated in General Social Surveys and in the 1975 Retirement Survey. Our results show that a significant portion of Canadians report enjoying life more after retirement compared to before retirement. Moreover, in 2002, three quarters of retired Canadians indicated being at least as satisfied with their finances as they were in the year prior to retirement. The most significant correlate of financial dissatisfaction that our analysis uncovers is involuntary retirement and, in particular, involuntary retirement associated with poor health.

The third chapter revisits a long standing question of whether households with higher *lifetime* income save a larger fraction of their income. The major difficulty in empirically assessing the relationship between lifetime incomes and saving rates is finding a credible proxy for lifetime income. Taking advantage of the unique characteristics of the Canadian Family Expenditure Survey data, we construct reliable lifetime income proxies. Our empirical analysis suggests that the estimated relationship between saving rates and lifetime incomes is sensitive to the instrument used to proxy lifetime income. Nevertheless, our preferred estimates indicate that, except for the

poorest households (who simply do not save), saving rates do not differ substantially across lifetime income groups.

The fourth chapter examines the effect of taxation on households' portfolio allocation decisions. The key challenge in empirically assessing this relation is to find a substantial and plausibly exogenous source of variation in marginal tax rates. We use variation in marginal tax rates across households with the same total earnings, which arise in progressive income tax systems with individual taxation. Employing the Canadian Survey of Household Finances, we find statistically significant but economically modest responses to differential taxation.

The fifth chapter investigates the problems associated with the estimation of intertemporal allocation parameters via linearized Euler Equations. We solve and simulate life-cycle consumption models in different economic environments and perform Monte Carlo experiments with these simulated data. Our results suggest that problems associated with the estimation of linearized Euler equations are strongly related to the assumed economic environments. In particular, the validity and relevance of conventional instruments used in the estimation depend on the concavity of the underlying policy rules, which in turn follows from features of the economic environment.

The final chapter reviews the economic literature on the saving behavior of lower income households. Our discussion focuses on the motives shaping the saving behavior of low-income households. In addition to the standard economic framework used for analysis of this kind, the *Life-Cycle/Permanent Income Hypothesis*, we also

discuss insights from "*behavioral*" economics. We review the international evidence on a variety of policies designed to stimulate savings among low income households.

Acknowledgement

I would like to thank my advisor Professor Thomas F. Crossley and the members of my thesis committee Professors Michael R. Veall and Paul Contoyannis for their invaluable assistance and mentoring during the course of my PhD studies. I would also like to acknowledge a great depth of gratitude to Professor Sule Alan both for her mentoring, and for the many ways in which she has been helpful to me.

I wish to thank my friends; Simo G. Goshev, Ulas Gunturkun, Mario Jametti, Logan McLeod, Mehmetcik Pamuk and Justin Smith for their friendship.

Lastly, but definitely not least, I thank my family for their unconditional support.

Preface

The second chapter, "The Adequacy of Retirement Savings: Subjective Survey Reports by Retired Canadians", was jointly authored by Professor Sule Alan and Professor Thomas F. Crossley, and has been published in Canadian Public Policy/Analyse de Politiques. I thank Canadian Public Policy/Analyse de Politiques for granting permission to reproduce it here. I contributed equally to all parts of the paper.

The third, fifth and sixth chapters were prepared jointly with Professor Sule Alan and Professor Thomas F. Crossley for journal publications. I contributed equally to all parts of these papers. The fourth chapter, "New Evidence on Taxes and Portfolio Choice", was prepared jointly with Professor Sule Alan, Professor Thomas F. Crossley and Dr. Sung-Hee Jeon and it has been submitted for journal publication. I was primarily responsible for the empirical analysis and participated in all other stages of the study.

Reference

Alan, Sule, Kadir Atalay, and Thomas F. Crossley, (2008) , "The Adequacy of Retirement Savings: Subjective Survey Reports by Retired Canadians," Canadian Public Policy , 34(1), 95-118.

Tables of Contents

List of Tables

List of Graphs

I

Introduction

This thesis is composed of independent chapters on household savings and finances. The next chapter investigates the living standards of Canadian retirees and the adequacy of their financial preparations for retirement. The third chapter examines differences in saving rates across lifetime income groups. The fourth chapter investigates the effect of taxation on households' portfolio allocation decisions. Problems associated with the estimation of linearized Euler equations are examined in the fifth chapter. The chapter before the conclusion reviews the economic literature on the saving behavior of lower income households. The last chapter concludes.

In recent years the adequacy of public and private provisions for retirement has been a hotly debated issue in many countries. The major difficulty when examining this issue is to find a way to compare household's pre- and post-retirement living standards. Conventional methods compare pre- and post-retirement levels of income, expenditure and consumption (see, for example, Banks, Blundell and Tanner 1998) to infer changes in the living standards. However, all of these measures face certain problems arising from the nature of retirement. For example, retirement is a period where income may be tighter but time is more plentiful. Evidence is beginning to accumulate that retirees are maintaining their living standards by substituting time for market expenditures. For example, they engage more in home production or spend more time on shopping for

better prices. Thus, retirees can achieve the same level of consumption with less market expenditure (for Canadian evidence, see Brzozowski and Lu, 2006).

In the chapter *"Adequacy of Retirement Savings: Subjective Survey Reports by Retired Canadians"* we employ an alternative, and more direct approach to assessing the living standards of retirees and the adequacy of their financial preparations. We examine retired Canadians' subjective survey reports of satisfaction with finances and with life, relative to the period before retirement. To the best of our knowledge, our paper is the first study that employs the subjective survey reports in this context in Canada.

Our results differ markedly from analyses of income and consumption. Firstly, contrary to hump shaped life paths of income and consumption, we find a flat profile for happiness over the life-cycle. Secondly, our results suggest that past cohorts of retiring Canadians reached that stage of life-cycle with adequate financial preparations. Finally, our results show that involuntary retirement, particularly because of ill-health, is a significant predictor of financial dissatisfaction.

The chapter *"Do the Rich Save More in Canada?"* examines the long standing question of whether households with higher lifetime incomes save a larger fraction of their income (i.e. whether they have higher saving *rates*). Although most non-economists find the proposition that the rich save more to be obvious, economists are more sceptical. An economist would point out that if agents are forward looking and try to smooth transitory income fluctuations, then a strong correlation between the saving rates and *current* income is expected, but this tells us little about the relationship between saving rates and *lifetime* incomes. Moreover, many widely used theoretical models of saving

behavior assume features that render saving rates that are *invariant* to lifetime income. As a result, the relationship between lifetime incomes and saving rates constitutes an important test of standard economic models of saving. However, finding a credible proxy for lifetime income poses a significant empirical challenge.

In this chapter we use the Canadian Family Expenditure Survey (FAMEX). The FAMEX provides us with excellent quality data on savings rates and also a number of potential instruments with which we can construct reliable lifetime income proxies.

This analysis is an important input into Canadian policy making, and a useful replication of a recent US study (Dynan et al. 2004) on data drawn from a similar economic environment. However, additional value is added by the unique features of the Canadian data. Our data enable us to exploit different saving measures and a variety of instruments that are not available in US data sets.

Our analysis suggests that the estimated relationship between saving rates and lifetime incomes is somewhat sensitive to the method used to construct a measure of lifetime income. Nevertheless, our preferred estimates indicate that, except for the poorest households (who simply do not save), saving rates do not differ substantially across lifetime income groups. These results support the use of standard economic models of saving when analyzing middle and high income households.

The chapter *"New Evidence on Taxes and Portfolio Choice"* investigates the effect of taxation on households' portfolio allocation decisions. Theoretical models predict that under a differential taxation system, households make portfolio allocation decisions based on not only the risk-return relationship of assets but also on their tax

characteristics. Given risk-return characteristics, households should minimize their tax burden by tilting their portfolios toward less heavily taxed assets. Moreover, less taxed households face a smaller incentive to invest in tax favored assets.

The key empirical challenge in testing these theoretical predictions is to find substantial and plausibly exogenous variation in tax rates. A number of recent studies (see for example, Poterba and Samwick 2003) try to get over this difficulty by studying changes in portfolio allocation around tax reforms (using a "*difference-in-difference*" approach). While this approach can identify successfully variations in the marginal tax rates (MTRs) *within* income groups, the results might be also very sensitive to the interval over which the data are *'differenced'*.

In this chapter, we estimate tax effects on portfolio allocation based on an alternative source of variation in tax rates. We exploit the variation in MTRs across households with the same total earnings, which arises in progressive income tax systems with *individual taxation* (such as the Canadian tax system). To the best of our knowledge, our strategy has not been employed in any other studies.

Using the Canadian Survey of Household Finances, we find statistically significant but economically modest responses to differential taxation. In a "placebo" test, using data from the U.S. (which has *joint taxation*), we also show that this is a true tax response and not attributable to heterogeneity in preferences or intra-household bargaining.

The chapter *"On the Identification of Intertemporal Allocation Parameters: Instruments and Sampling Issues"* investigates the problems associated with the

estimation of intertemporal allocation parameters via linearized Euler equations. There exists a large and growing empirical literature that examines the sensitivity of consumption (or equivalently saving) to interest rates (the elasticity of intertemporal substitution) and to uncertainty (prudence). Most empirical work estimates these effects by employing an optimality condition (Euler equation) derived from the dynamic optimization problem of a consumer. Although this approach has several advantages, there are also problems associated with it. The main problem is that the non-linear structure of the Euler equations precludes their estimation from data that contain measurement error. In the literature, this problem is addressed by linearizing the Euler equation and then using standard linear IV techniques to deal with the effects of measurement error. However, it has recently been argued that the approximation bias induced by linearization may be worse than the problems that linearization is intended to solve. Consequently, the usefulness of the Euler equation approach has been called into question (eg., Carroll, 2001 , Attanasio and Low, 2004).

In this chapter, we explore this issue with a series of simulation studies. We show that: (1) the problems associated with linearized Euler equations are strongly related to the assumed economic environment. In particular, the degree of approximation bias is related to the concavity of the underlying policy rules, which in turn follows from parameters of the economic environment; (2) there exist environments in which the approximate Euler equation can yield unbiased estimates (as well as environments in which the approximate Euler equation works quite badly); and, (3) careful sample

selection (or conditioning) can lead to better estimates of intertemporal allocation parameters in a range of economic environments.

The final chapter "*Increasing the Saving Rates of Low Income Families*" reviews the economic literature on the saving behavior of lower income households. The last quarter century has witnessed the introduction of numerous government programs intended to stimulate saving (tax-favored saving accounts, employer/government matching contribution and bonus). However, there are still many households that save very little. Several explanations for this failure of policy have been proposed in the literature. One explanation is that the means testing in social insurance programs provides a strong disincentive for saving by low income households. Alternative explanations include the necessity of an absolute minimum consumption level, which precludes saving at low income levels, or a limited capacity to plan among low income households, or lack of financial literacy. Since each of these alternative explanations has very different policy implications, it is important to understand the role of each alternative in shaping saving behavior of low-income households.

In this chapter, we show that the standard economic framework, the *Life-Cycle/Permanent Income Hypotheses*, as well as emerging approaches coming from "*behavioral*" economics offer valuable insights for understanding the saving behavior of low income households. We also review evidence on the efficacy of a variety government programs designed to stimulate saving among this low income households.

References

Attanasio, O. P. and H. Low, (2004), "Estimating Euler Equations," Review of Economic Dynamics, 7(2), 405-435

Banks, J., R. Blundell, and S. Tanner, (1998), "Is there a retirement-savings puzzle?" American Economic Review, 88(4), 769–88

Brzozowski, M.and Y. Lu, (2006), "Home Cooking, Food Consumption and Food Production among the Unemployed and Retired Households," SEDAP Working Paper, no:151.

Carroll, C., (2001), "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)", Advances in Macroeconomics, 1(1), 1003-1003

Dynan ,K. E., J. Skinner, and S. P. Zeldes, (2004), "Do the Rich Save More?" Journal of Political Economy, 112, 397-444

Poterba, J. and S. Andrew, (2003), "Taxation and household portfolio composition: US evidence from the 1980s and 1990s," Journal of Public Economic, 87, 5-38.

II

The Adequacy of Retirement Savings: Subjective Survey Reports by Retired Canadians[1]

1. Introduction

Does public policy, in conjunction with Canadians' own preparations, result in Canadians having adequate financial resources in their retirement years? Has the adequacy of retirement provisions changed over time? These are obviously important questions - and they can only become more important as the population ages. However, they are also very difficult questions to answer.

One obstacle that must be overcome in answering these questions is to define a notion of adequacy. Like most economists, we would suggest that retirement preparations are appropriate if the individual or household can't be made better off (in an ex ante sense) by reserving additional financial resources from earlier life stages to use in their retirement years, or by bringing some resources forward from retirement to an earlier stage of life. We also accept that, in general, individuals prefer constant living standards to fluctuations. Together, these ideas imply (roughly) that appropriate retirement provisions equalize living standards between pre- and post-retirement years. This captures the essence of the life cycle framework for modeling consumption and saving behaviour, which is favoured by many economists (see, for example, Browning and

[1] This chapter was published in volume 34 of Canadian Public Policy. Full reference is "Alan, Sule, Kadir Atalay, and Thomas F. Crossley, (2008) , "The Adequacy of Retirement Savings: Subjective Survey Reports by Retired Canadians," Canadian Public Policy , 34(1), 95-118."

Crossley, 2001).[2] It also accords well with common sense: many Canadians would consider adequate financial preparations to be those that allow individuals to enjoy a standard of living in retirement years similar to that which they became accustomed to during their working life.

It is important to distinguish adequate *savings* from adequate *resources*. An individual could have adequate savings in the sense just defined, but still have resources in retirement that society would deem inadequate. That is, an individual or household might have allocated their lifetime resources across pre- and post-retirement periods as best they can, but still experience deprivation in retirement (and before retirement) because of a lack of lifetime resources. This should not be interpreted as failure of private saving behaviour or public pension policy.[3] Deprivation that arises in this way is appropriately addressed by policies targeted at lifetime earnings capacities and general inequality.

How then can we judge whether public and private provisions for retirement are sufficient to maintain pre-retirement living standards into post-retirement life? One practical idea is to examine incomes. The problem with this, of course, is that income is only one component of financial resources. If Canadians enter retirement with non-annuitized wealth, and dis-save that wealth in retirement, then income will understate their financial resources, and hence standard of living, in retirement. For this reason, many economists prefer to examine expenditure (or consumption), pre- and post-

[2] More formally, life cycle models of optimal saving behaviour imply that forward looking households should allocate resources between current consumption and future consumption (saving) so as to equalize the discounted expected marginal utility of expenditure.

[3] Indeed, by this definition adequate savings for some low income individuals may be no savings at all (because they will have access to means tested public benefits).

retirement (see for example, Banks, Blundell and Tanner, 1998). The idea is that what people spend is a good indicator of the financial resources at their disposal. Moreover, the thinking goes, total expenditure is closely related to total consumption, which is a natural measure of material well-being – and therefore expenditure is closer than income to what should be "smoothed" (or equalized between the pre- and post retirement periods of life).

However, even total expenditure is an incomplete measure of living standards, and the relationship between expenditure and material well-being may differ between the pre- and post-retirement years. First, many goods have a "public" nature to them so that it is uncontroversial that larger households enjoy some economies of scale ("two can live cheaper than one"). Individuals live in households of different sizes, and in particular, Canadians tend to spend their retirement years in smaller households (after their children have moved out, for example.) Thus comparisons of pre- and post-retirement expenditure levels will depend critically on adjustments made for the different needs of households of different sizes. While such adjustments are in principle easy to make, it is rather more difficult to know whether one has made the right adjustment. (Does a household of two persons need half the expenditure of a household of four persons to have the same standard of living? Or do they require two thirds of the expenditure? Or three quarters?)

Second, retirement is a period when income may be tighter, but time is more plentiful. Evidence is beginning to accumulate that retired households maintain material living standards by substituting time for market expenditures. They can do this in a number of ways. For example, they may shop more in order to take advantages of sales and other price variability, and hence lower the expenditure required for a constant

quantity of goods and services. Additionally, they may engage in more home production (substituting ingredients and cooking time for the purchase of prepared foods, for example.) For further details, and compelling evidence, see Aguiar and Hurst (2005a, b) and Brzozowski and Lu (2006) (of which the latter analyses Canadian data.) The implication is that assessing living standards with consumption data requires a careful modelling of both spending and time use, which is a daunting task.

Given these difficulties, it would be useful to have another way to assess the living standards of retired Canadians, and the adequacy of their financial preparations. In this paper we examine retired Canadians' answers to subjective questions regarding their satisfaction with their current financial circumstances and with their life in general. Although economists have traditionally focused on quantitative measures of behaviour (such as incomes and expenditures), the use of subjective survey measures of economic well-being has been rising in recent years (see Frey and Stutzer (2002) for a survey).

It turns out that Statistics Canada has been asking such questions of retired Canadians, in a variety of surveys, stretching back quite a few years. Moreover, some of the relevant satisfaction questions are posed *relative to* the respondents' pre-retirement years. As such, they line up nicely with our notion of adequacy, which hinges on the equality of pre- and post- retirement living standards. For example, both the 1994 and 2002 General Social Surveys ask:

> *Compared to the year before you retired, would you now say that you are better off financially, worse off or about the same?*

If the models in the life cycle framework are broadly correct, then retired Canadian seniors should on average report that their financial circumstances are about the

same as before retirement. Of course some households may experience positive shocks, and some households may experience negative shocks, but financial circumstances should not be substantially worse (or better) after retirement *on average.* If a majority of retired Canadians indicated that they experienced living standards in retirement that were higher than prior to retirement, that might indicate excessive saving (where saving is broadly defined, and includes, for example, the accrual of CPP/QPP entitlements). On the other hand, if a majority of retired Canadians indicated that they experienced retirement living standards that were lower than prior to retirement, this would indicate that financial preparations (savings) for retirement were not adequate (again, on average).

A more general question about enjoyment of life, again *relative* to before retirement was asked in three General Social Surveys (1989, 1994 and 2002), and we examine this as well.

The objective of this short paper is to assess what the responses to these questions indicate about the adequacy of financial preparations for retirement among retired Canadians, past and present. The novelty of the analysis lies in the use of subjective survey responses to assess living standards, rather than data on income or expenditure.[4]

[4] We are aware of one other paper to take this approach with Canadian data. Baker et al. (2005) use cross cohort variation in entitlements to examine the effect of government retirement income programs on the well being of the Canadian retirees. In addition to income and consumption, they use self-reports of happiness in the General Social Surveys as a well being measure. Their analysis differs from ours in that they focus on the effect of policy changes on the well-being of seniors, rather than on a comparison of pre- and post-retirement living standards (which is our focus). They naturally therefore focus on the direct happiness questions in the GSS (which we also examine in Section 2) but do not employ the *relative* satisfaction questions that are our main focus.

These subjective data are not necessarily superior to those more quantitative measures, but they suffer from different limitations and so are certainly complementary.

In the next section, we consider the life cycle path of "happiness", and contrast it with life cycle paths of income and expenditure. This provides further motivation for our main analysis, which is of the "relative satisfaction" questions described above, and which is reported in Section 3. Section 4 provides a concluding discussion.

2. Age-Profiles of Income, Expenditure and Happiness

To further motivate what follows, we first look at broad patterns of income, consumption and happiness over the life cycle. These are presented in Figure 1. The top right hand panel presents the age profile of average real equivalent disposable income: household income, net of taxes and transfers, has been deflated to 1992 Canadian Dollars using the Current Price Index, and divided the commonly used 'root of household size' equivalence scale. The top right panel presents the age profile (or life cycle path) of average real equivalent expenditure on nondurable consumption. Finally, the bottom left hand-side panel presents the age-profile of happiness, where happiness is measured as the % of respondents reporting they are "very happy" or "somewhat happy".[5] The income and consumption figures are based on the 1992 Family Expenditure Survey and the 1998 Survey of Household Spending. The happiness figure is based on the 1990 and 1998 General Social Surveys.[6] In each case we use two surveys because it is well known that

[5] The exact question is: *Presently, would you describe yourself as...very happy, somewhat happy, somewhat unhappy, very unhappy?*

[6] General Social Surveys in a number of other years collected the happiness reports that we use to produce this picture. However, these two particular surveys were chosen because the public use file of the GSS for these two years had a continuous (rather than categorical) age measure. The continuous age measure allows

cross-sectional age profiles confound both age and cohort effects.[7] In these figures, the solid lines connect observations on the same five-year birth cohort, and the number associated with each dot gives the first year of the five-year birth cohort.[8] While some cohort differences (at a given age) are apparent, the overall shape of the age profiles is quite clear.

The income and consumption figures exhibit a definite hump shape. Using income or consumption as a well-being measure suggests that well-being rises through early life, peaks in middle age, and falls thereafter, with a substantial decline in normal retirement years. The patterns shown here replicate very well known results, and they are not sensitive to the data we use, or to various measurement choices.[9]

In stark contrast, the age-profile of happiness is quite flat. The shape of this profile is robust to our treatment of the categorical responses.[10] While it is well known that happiness rises with income within a group of respondents at a given point of time, this figure illustrates that average happiness does not seem to track the movements of income over the life cycle.[11] Self-reports of happiness seem to tell a quite different story

---

us to line up birth cohorts in successive surveys exactly. The FAMEX/SHS surveys were chosen to match the timing of the GSS surveys as closely as possible. There was a FAMEX survey in 1990 but it is only representative of Canadians in major urban centres. The surveys we employ to produce Figure 1 sampled both urban and rural Canada.

[7] Within a given year, older individuals are also born earlier. If older cohorts have lower lifetime earnings (because of productivity growth, for example), then their entire age profile may be lower than that of latter cohorts. When individuals of different cohorts are combined in a cross-sectional age profile, the lower incomes of earlier cohorts may be misinterpreted as a decline with age.

[8] So, for example, the cohort born in 1920 through 1924 is labelled "20".

[9] See Browning and Crossley (2001) for U.K. evidence and further references to the international literature. For Canadian evidence see Robb and Burbidge (1989) and, more recently, Crossley and Pendakur (2006).

[10] In fact, the fraction of respondents reporting each of the four possible responses is largely invariant to age. Thus any combination of categories or cardinal treatment of the responses gives a similar profile.

[11] This is somewhat reminiscent of the now famous finding that average happiness in a country does not rise over time with improvements in per capita income (Easterlin 1974, 1995).

about the evolution of well-being with age, and about the well-being of Canadians of normal retirement age.[12]

Figures such as these infer the evolution of average well-being with age from the responses of (distinct) samples from a given birth cohort at different points in time. In our main analysis, which is reported in the next section, we instead use question about satisfaction relative to the year before retirement to infer the changes in living standards experienced by individuals.

3. Relative Satisfaction with Life and Finances

3.1 Data and samples

The data we employ in this section are drawn from four Statistics Canada surveys, which together cover a quarter of a century. They are:

(i) Cycle 16 of the General Social Survey (2002).

(ii) Cycle 9 of the General Social Survey (1994),

(iii) Cycle 5 of the General Social Survey (1989),

(iv) The 1975 Retirement Survey

Comparisons over time rest on the comparability of survey design and coverage, of survey questions, and of general economic conditions. With respect to the latter, it is obviously the case that with only four surveys, it is difficult to discern permanent changes (having to do with changes in policy and behaviour) from more cycle or transitory developments (such as episode of significant job loss, or high inflation). Figure A1 in the appendix presents unemployment and inflation rates for Canada for the period spanned

[12] Some international evidence suggests a "u-shaped" pattern of happiness over the life cycle (for example, Blanchflower and Oswald, 2007). However, in Canadian data we find neither a "u-shape" nor a "hump".

by these surveys. In 1975 inflation was high, and unemployment was moderate. The 1989 survey comes at the end of the long expansion of the late 1980s: inflation was moderate and unemployment was at its lowest point in a decade. Relative to 1975, inflation is lower and unemployment comparable. In 1994 and 2002 inflation is lower still. Unemployment is higher in 1994 than in 1989 or 1975 as the economy was moving out of the recession of the early 1990s. In 2002 unemployment was roughly comparable to the levels seen in 1975 and 1989, even though it had risen after the end of the late 1990s boom.

With respect to survey design and content, the three General Social Surveys are similar in their design and coverage, and ask common questions. Hence data from these three surveys are broadly comparable. Comparisons with the 1975 Retirement Survey are more difficult. Different questions are asked, and the design and coverage of the survey is also different. There is nothing we can do about the former. The only thing we can do about the latter is to find broadly comparable subsets of the data. For example, the 1975 Retirement Survey surveyed only individuals over 55 years of age, and it is obviously not difficult to impose this sample restriction on the data from other years. Overall, we believe that comparisons over time are of interest, even if they are imperfect, so long as we are mindful of the inherent limitations.

Turning to the details of the individual surveys, the *1975 Retirement Survey* was conducted in February of 1975. The sample was drawn from households rotating out of the Labour Force Survey. If a respondent was over 55, a survey was left for self-enumeration. Screening questions determined whether the respondent received the

retirement survey or a pre-retirement survey.[13] There were 1590 responses, corresponding to a response rate of approximately 80%. From these we deleted 54 observations with missing age information and 141 observations whose Labour Force Survey responses indicated they were working or looking for work. These restrictions resulted in a working sample of 1395 observations.

The *General Social Survey* series was introduced in 1985 and continuously conducted each year since. The objective of GSS is to monitor changes in the well being of Canadians over time and provide immediate information on specific social issues of current interest. Our analysis is based on cycles 4, 9 and 16. The response rates of these surveys were respectively 81, 81 and 86%, which compares well with other Canadian household surveys.[14] Further information about these surveys is summarized in Appendix Table A1.

For cycles 4 and 9 the target population is Canadians 15 years of age older while in Cycle 16 the target population was restricted to persons aged 45 and over.

In the *1989 GSS* (cycle 4) the core content concerns education and work. Module K of the survey includes questions about respondents' retirement experience. Of the 9338 valid respondents to the survey, 1143 were screened into this retirement module. From these we kept only respondents who declared their main activity during the previous year (1988) as "retired" and deleted those that responded "no" to the question

[13] The screening questions were: "Are you working?", "Are you looking for work?" and "Do you consider yourself permanently retired?" Entry into the Retirement Survey required the sequence of answers "No", "No", "Yes".

[14] For example, the Statistics Canada's household expenditure surveys (two of which were used to produce the income and expenditure numbers underlying Figure 1) typically achieve response rates around 75%.

"have you ever worked at a job or business?" In order to be consistent with the 1975 Retirement Survey, we kept only respondents aged 55 and older. These selections resulted in a final sample of 1003 respondents.

The core content of the *1994 GSS* (Cycle 9) concerns the transition into retirement and post-retirement activities. Of the 11,876 valid responses we again selected those that were screened into "Section K" on retirement, currently self report being retired, and who used to work at a job before retirement, and who are aged 55 or older. This left us with a working sample of 1537 respondents.

The *2002 GSS* (Cycle 16) was designed to provide data on the aging population, and contains information about person's retirement preferences and experiences. There are 24,870 valid respondents, all aged 45 years or over. The GSS first asks the main activity of the people in last 12 months. If the respondent declared this as "retired", she or he is defined as a retiree. The rest are asked a second screening question: "Have you ever retired?" If the answer of this question is yes they are also classified as a retiree. From this pool we selected respondents who are aged 55 years or over; who are not looking for work; and have not worked since their retirement. This left a working sample of 7940 retirees.

Our analysis is based on publicly available micro-data files, and in all calculations we use the sample weights provided in these files by Statistics Canada.[15] Descriptive statistics for our sample from the 2002 GSS are provided in Table A2.

[15] One consequence of using household survey data is that institutionalized individuals are typically outside of the frame of the survey. Thus the group of retired individuals with particularly poor health and particular financial needs are missed. This problem is not specific to this paper but plagues any research on the older population that employs household survey data – including income and expenditure surveys.

3.2 Satisfaction measures

These are cross-sectional (rather than panel) surveys. However, it is the retrospective nature of the financial satisfaction questions in the 1994 and 2002 General Social Surveys (and the general satisfaction questions in 1989, 1994 and 2002) that allows us to study a life cycle (that is, dynamic) question. In particular, as noted above, the satisfaction questions ask about satisfaction *relative* to the year before retirement:

Compared to the year before you retired, would you now say that you are better off financially, worse off or about the same?

Compared to the year before you retired, do you now enjoy life more, less, or about the same?

Thus, if these questions are answered properly, they should capture *changes* in financial satisfaction or living standards, across the retirement event.

An obvious methodological concern is the validity of the subjective satisfaction measures. There is currently considerable disagreement among social scientists, and even among economists, as to the value of such measures. One can compare, for example, the quite negative view of Bertrand and Mullainathan (2001) with the more positive stance of Frey and Stutzer (2002) or Kahneman and Krueger (2006). A key issue with subjective questions is the comparability of responses across individuals or groups of individuals (whether individuals interpret the scale in the same way.) A second advantage of the fact that the questions we study in this section are asked *relative* to an earlier time period is that a relative formulation diminishes concerns about interpersonal comparability somewhat. Simply put, the responses "better off, about the same, worse off" seem less subject to individual interpretation than an absolute scale ("unsatisfactory, somewhat

satisfactory, very satisfactory"). More formally, if individual assessments of financial or life satisfaction are comparable up to location, then asking a question about changes may effectively eliminate ("difference-out") remaining heterogeneity in response behaviour.[16,17]

Although the questions are retrospective in nature, at best they provide (relative) information about two points: current circumstances and the year just before retirement. They are potentially useful for determining whether Canadians' (self-assessed) wellbeing changes at the point of retirement. As income and expenditure fall at retirement, and those changes are often interpreted as a fall in living standards, the ability to assess changes in satisfaction right around retirement is valuable. On the other hand, with two points, it is obviously not possible to rule out more complicated dynamic patterns of satisfaction.[18]

Another concern with the subjective responses is the possibility of "adaptation", a phenomenon well known in the literature (Kahneman and Krueger, 2006). A crude strategy for assessing whether adaptation is a serious problem in our analysis is to repeat the analysis on the subset of individuals who are quite recently retired at the time of the survey.

---

[16] This would be true if response-scale heterogeneity could be modelled as an additive individual "fixed-effect".

[17] The current best practice for assessing the comparability of subjective responses across individuals or groups involves the use of "vignettes" (see, for example, King, Murray, Salmon and Tandon (2004) or Kapteyn, Smith and van Soest, (2006)). In this paper we report analysis of the best currently available data, but a natural future extension would be to collect new data using vignettes or other methods to assess the comparability of subjective responses.

[18] Specifically, a scenario that these questions would not detect is this: people are not forward looking enough to make preparations for retirement, but are sufficiently forward looking enough to regret this before they actually arrive in the year prior retirement.

Finally, it may be that these questions are easier to answer when retirement is a sharp event, rather than gradual.

3.3 Levels of Relative Satisfaction with Life and Finances

We begin, in Table 1, by tabulating responses to questions about satisfaction with life. Columns 1, 3, 4 and 5 report calculations based on the RS75 and 1989, 1994 and 2002 GSS, respectively. In each case, the calculations are based on our standard samples of retired Canadians aged 55 years and over. The GSS question is as described above. The most comparable question in the RS75 asks *"Is any aspect of life better or worse than expected?"* This question shares the relative aspect of the GSS question although it is relative to pre-retirement expectation, rather than relative to pre-retirement experience. It is also important to note that having an aspect of life better than expected, and having an aspect of life worse than expected are *not* mutually exclusive possibilities (in contrast, with the GSS question, a respondent cannot be *both* more and less satisfied.) As previously noted, these question differences must be borne in mind when making comparisons between the surveys.

In 1975 about equal numbers (23 and 25 percent) of retired Canadians indicated that some aspect of retirement life was, respectively, worse, and better, than they had expected.

In the GSS surveys, about 40 percent of Canadians report that they enjoy life about the same as in the year prior to retirement (39, 34 and 46 percent in 1989, 1994 and 2002, respectively.) More than 40 percent report that they enjoy life *more* than prior to retirement (44, 48 and 42 percent in 1989, 1994 and 2002). The number of retired

Canadians, 55 years of age and over, that report enjoying life less than in the year before retirement never exceeds 20 percent, and only 12 percent in 2002.

The rest of Table 1 reports the same calculations for two different subsets of the data. The first sample limitation we consider, in columns 2, 6, 7 and 8, is to focus specifically on retired Canadians aged 70 and older. There are two reasons to do this. First our samples are of retired Canadians. If only those Canadians that can afford retire do so, then our sampling is endogenous to the outcome we are examining, and the numbers we report could be misleading. Second, patterns of retirement have changed over time, and perhaps even the notion of what it means to be retired has evolved. This means that the process of selection into our sample may differ across years, which would compromise any comparisons over time.

At any age, the higher the fraction of the population that is retired, the less scope there is for selection issues to influence our results.[19] As we move from samples of retired Canadians aged 55 and above to samples of retired Canadians aged 70 and above, retirees become a larger fraction of Canadians in the relevant age group. Thus the potential selection affects should be attenuated. It is certainly *not* our contention that our calculations on samples of retired Canadians aged 70 and above are immune to selection problems. Rather, our hope is that significant selection problems might reveal themselves in the comparison of the two sets of numbers (columns 1, 3, 4, 5 against 2, 6, 7 and 8). The limitation of this kind of robustness check, of course, is that differences between the two sets of numbers could also reflect genuine age, cohort or time-since-retirement

[19] In the extreme, if all Canadians above an age k are retired in all years, then tabulations of responses from Canadians above age k would not suffer from either of the selection problems outlined above.

effects, rather than selection.

The result, as Table 1 reveals, is that limiting the analysis to those 70 years of age and older has little effect on the time path of satisfaction levels. For example, in either sample, the fraction of retirees indicating that they enjoy life less (than in the year prior to retirement) fell from1989 to 1994, and then fell further to 2002. This robustness of the time patterns in the data gives us some hope that they are not simply the consequence of the changing nature or conception of retirement.

On the other hand it is true that, in every year, the number of retirees stating that they enjoy life less than the year before retirement is higher in the older sample. In each year, the fraction stating that they enjoy life more is correspondingly lower in the older sample. It is possible that this difference reflects more selection in the younger sample (that is, at younger ages, those who are not financial prepared do not retire, but at older ages there is less scope for continued work). Alternatively, it may reflect genuine age effects (arising, for example from declining health), or even cohort effects.

A second robustness check is reported in Column 9. Here, for the 2002 GSS only, we tabulate responses to the relative life satisfaction question only for those respondents within 5 years of their first retirement. Because the satisfaction question asks the respondents to make a comparison to the year prior to retirement, focusing on this group captures comparisons made over a shorter time span. One reason to do this is the concern that the self-reports suffer from adaptation, whereby individuals finding themselves in less advantageous circumstances eventually adjust their expectations to their new circumstances and return to previous levels of self-reported well-being. Comparing

Column 9 to Column 5 indicates that in 2002 comparisons of life to the year before retirement were a bit more favourable among recent retirees than among retirees overall. Just as with our first robustness check, this comparison potentially confounds multiple effects (for example, an age effect with an adaptation or other time-since-retirement effect). Nevertheless, the general stability of the response distribution across columns suggests to us that the broad picture given by these numbers is correct.

In RS75 respondents were asked to identify which aspects of life were better or worse than expected (if any). In the 1989 and 1994 GSS (but not 2002) respondents who indicated that they enjoyed life less or more than prior to retirement were asked a follow up question about the reasons for this.[20] A partial tabulation of the responses is given in Table 2. The categories of response we report are income/economic and health. As in Table 1, we report the distribution of responses first among retired Canadians aged 55 and over (Columns 1, 3 and 4) and then among retired Canadians aged 70 and over (Columns 2, 5 and 6). In each cell there are two numbers. The top number indicates the fraction of those reporting that they enjoyed life less that gave this reason (the conditional probability of giving this response.) The lower number (in square brackets) gives the fraction of the entire sample that said they enjoyed life less *and* gave this response as the reason (the unconditional probability of being dissatisfied for this reason).

The number reported in Table 2 indicate that in 1989 and 1994, when retired Canadians report that they enjoy life less than prior to retirement, they are much more likely to attribute their dissatisfaction to health problems than to economic reasons or

[20] They were asked: *What is the main reason that you now enjoy life less (more)? Is it....*

income. In 1994, of the 17 percent of retirees that said they enjoyed life less than prior to retirement only 12 percent cited income or economic circumstances (Table 2, Column 4) as the reason. Combining these numbers, only 2 percent of retirees in 1994 indicated that they were enjoying life less than before retirement because of income or economic reasons. Four times as many cited health as the reason for enjoying life less than before retirement.

In Table 2 a significant difference is apparent between RS75 on the one hand, and the 1989 and 1994 GSS on the other. In RS75, more retirees cited income/economic circumstances than health as an aspect of life that was worse than expected. The change in the distribution of reasons for discontent between 1975 and 1989 might reflect the significant increase in public transfers to the elderly in the 1970s and the consequent reduction in poverty among seniors that has been well documented. However, comparisons between the RS75 and the 1989 and 1994 GSS should be made with considerable caution: the question posed to respondents is undeniably different.

Table 3 follows the same format as Table 1, except that it summarizes responses to the question about relative *financial* satisfaction. This question was not asked in the 1989 GSS and there is no comparable question in the RS75.[21]

In 2002, 74 percent of retired Canadians reported that they were better off financially (18 percent) or about the same (56 percent) compared to the year before they retired (Table 3, Column 2). The numbers for 1994 (Column 1) are similar. When we restrict the sample to retirees aged 70 and older, the distribution of responses is slightly

[21] The RS75 does contain questions about income satisfaction, but this question is not asked relative to pre-retirement expectation or experience.

more positive (Columns 3 and 4). On the other hand, when we focus on those that retired within the last 5 years (Column 5) we see a slightly higher fraction of negative responses (compare 31 percent with 26 percent in Column 2).

Overall, these numbers are broadly compatible with the view that in the recent past, the combination of public provisions and private preparations has delivered adequate financial preparation for retirement, at least in an *ex ante* sense. In both years (1994 and 2002) and in all the samples we consider, more than half of retirees describe themselves as being neither better nor worse off financially, compared to before retirement. These households claim to have successfully smoothed living standards. Of course, some households do report some deterioration of their financial position, but in most samples a fairly comparable fraction report being *better* off financially in retirement than before retirement. Against a standard of constant living standards, this latter group may have *over* saved.

3.4 Correlates of Relative Satisfaction

Table 1 through 3 reveals important heterogeneity. Some retired Canadians report being less satisfied with their finances than prior to retirement, while others report greater financial satisfaction than prior to retirement, and similarly for overall satisfaction with life. In Tables 4, 5 and 6, we relate these outcomes to observable characteristics of respondents and their families. The analysis is of course descriptive, and great caution should be exercised in ascribing causal interpretations to the associations in the data.

In Table 4 we tabulate the distributions of responses to the `relative' satisfaction questions conditional on values of covariates (where we use `relative' to emphasize that

satisfaction is measured relative to the year before retirement.) The analysis is bivariate: we consider one covariate at a time. A number of interesting associations are immediately apparent. Married individuals, of either gender, report greater relative life satisfaction than those living alone, but there is little difference in the relative financial satisfaction of the two groups. Relative life satisfaction appears to decline with age while relative financial satisfaction rises. Again, age may be capture the effects of age, of time since retirement, of birth cohort, or of differential selection. Relative life satisfaction is strongly positively associated with education, but relative financial satisfaction is not. Among retirement reasons, those who retired because of unemployment are most likely to report that they enjoy life less than prior to retirement, and that they are worse off financially.

Home ownership and receiving a pension benefit from a former employer are also positively associated with relative life satisfaction. Relative financial satisfaction is also higher among retirees who receive pension benefits from a former employer.

In Tables 5 and 6 we turn to multivariate analyses and report estimates of Ordered Probit models for relative life and financial satisfaction respectively. The models are estimated standard maximum likelihood methods. Ordered Probit models are appropriate for modelling categorical variables in which the response categories can be ranked, as is the case with the satisfaction scales considered here. The dependent variable takes the value 0 if the respondent enjoys life less (is worse off financially), 1 if the respondent is about the same, and 2 if the respondent enjoys life more (is better off financially). The use of ordered Probit models with satisfaction or happiness data is quite standard (see for example, Ferrer-i-Carbonell and Frijters, 2004).

In addition to estimated coefficients, we report marginal effects on the combined probability of enjoying more and about the same (being better off or about the same). Of course, this is the negative of the marginal effect on the probability of enjoying life less (being worse off).

For both outcomes we estimate two models. The first specification (on the left) conditions on sets of dummy variables capturing age, education, province and family type as well as a dummy variable for home ownership, a dummy variable indicating receipt of a pension from a former employer, and a dummy indicating that retirement was "voluntary". In the second specification (on the right) the voluntary retirement dummy is replaced by a set of dummies capturing alternative retirement reasons in greater detail.

The models for relative life satisfaction (Table 5) indicate a number of statistically and economically significant partial effects. First, there is a strong negative age profile in relative life satisfaction. Females report greater relative satisfaction with life than males. Retirees are more likely to report that they enjoy life more than prior to retirement if they own their own homes, or if they receive a pension benefit from a former employer. Finally, there is a strong positive association between relative life satisfaction and voluntary retirement. Retirements because of poor health or caretaking (presumably a partner's poor health) or because of unemployment or business closure are all very strongly associated with enjoying life less than before retirement.

Turning to relative financial satisfaction, we again see a strong age profile, but in the opposite direction: relative financial satisfaction rises with age. As discussed above, this might reflect a genuine age effect, a cohort effect, or perhaps an adaptation or other

time-since-retirement effect. It might also reflect a kind of selection effect, generated by differential mortality (with the financially secure outliving their less fortunate peers.)

The only other strong association with relative financial satisfaction is voluntary retirement. Respondents who report that they retired voluntarily are 11 percentage points less likely to report that they are financially worse off than before they retired. When retirement reasons are broken down further, we uncover a very strong negative association between retirement because of poor health and relative financial satisfaction.

4. Summary and Conclusions

In this short paper we have explored response by Canadian retirees to subjective survey questions administered in General Social Surveys and in the 1975 Retirement Survey. We have documented a number of potentially important facts. They are:

i) Happiness does not track income over the life cycle. On average, income is hump-shaped, as is consumption. Average happiness is flat across age, including normal retirement years.

ii) Many more retired Canadians report enjoying life more than before retirement than the converse; more than 80% enjoy life more or the same as the year prior to retirement. This was true in 1989 and 1994, as well as in the most recent available data, from 2002.

iii) Retired Canadians who report that they do not enjoy life as much as in the year prior to retirement are much more likely to cite health concerns than economic reasons for their dissatisfaction.

iv) In 2002, three quarters of retired Canadians reported being either as satisfied or more satisfied with their finances than they were in the year prior to retirement. Almost as many reported an improvement in financial situation as reported a decline.

v) Life satisfaction appears to fall with age while financial satisfaction appears to rise with age (both measured relative to the year prior to retirement). An important caveat here is that the cross-sectional data available to us may confound age, selection and cohort effects.

Subjective survey reports of relative financial or life satisfaction seem to paint quite a different picture than analyses of income or expenditure. On their face, these survey responses would seem to suggest that past cohorts of retiring Canadians reached that stage of the life cycle with adequate financial preparations. Again, this is in the sense of not having disruptions in living standards around retirement.

Recent U.S. studies (Ameriks et al., 2007; Hurd and Rohwedder, 2006) indicate that while there is a decline in expenditures at retirement, this decline is fully anticipated by households, and as such, probably does not represent a fall in living standards. These

anticipated expenditure declines appear to be associated with declining consumption needs and with the substitution of time inputs for market expenditures. Comparable expectations data are not currently available for Canada. If the average expenditure falls around retirement apparent in Canadian data (as in Figure 1) were similarly anticipated, and related to the same factors, it would explain why most retired Canadians do not report a decline in their financial situation, and many report an improvement.

An additional important point is that some retired Canadians *do* report deterioration in their financial circumstances. Because no more report deterioration than improvement, this does not indicate a systematic failure of preparation. It may be that there is a subset of Canadian households for whom the combination of private actions and public provisions does not deliver adequate financial resources for retirement (and, equally, another subset that *over*-saves). Alternatively, it could be that as they approach retirement, Canadian households face risks that neither private nor public arrangements fully insure. Most notions of adequate ex ante preparations still allow the possibility of negative (and positive) surprises, ex post. The last important fact produced by our analysis is consistent with this second hypothesis. It is:

vi) The most significant correlate of relative financial dissatisfaction that our analysis uncovers is involuntary retirement, and in particular, involuntary retirement associated with poor health.

Again, suggestive connections can be made to the international literature. For example, in a recent study of expenditure falls around retirement that employs panel data from the United Kingdom, Smith (2006) shows that spending falls only among (the households of)

men who retire involuntarily. Findings such as these suggest areas where policy innovation might be most fruitfully targeted.

There are many problems with subjective survey self-reports of financial satisfaction. Social scientists and policy makers are right to be cautious, and even sceptical when confronted with such data. Nevertheless, our view is that, at a minimum, the data summarized in this paper suggest that income and expenditure data also be interpreted carefully.

Further analysis and data are required. Most helpful would be a Canadian panel survey that collected incomes, expenditures and a range of other living standard measures (including subjective ones) around retirement. Data collection exercises matching this description are underway in many other advanced countries.

References

Aguiar, Mark, and Erik Hurst, (2005a) "Consumption versus Expenditure," *Journal of Political Economy,* 113, 919-948

Aguiar, Mark, and Erik Hurst (2005b) "Life cycle prices and production," paper presented at the 2005 National Bureau of Economic Research Summer Institute, Cambridge MA., July 2005.

Ameriks, John, Andrew Caplin and John Leahy, (2007), "Retirement Consumption: Insights from a Survey," *Review of Economics and Statistics*, 89(2):256-274.

Baker, Michael, Jonathan Gruber and Kevin Milligan, (2005), "Retirement Income Security and well Being in Canada" manuscript, University of British Columbia, December 2005.

Banks, James, Richard Blundell and Sarah Tanner, (1998), "Is There a Retirement-Savings Puzzle?", *American Economic Review*, 88(4):769-788.

Bertrand, Marianne and Sendhil Mullainathan (2001) "Do People Mean What They Say? Implications for Subjective Survey Data." *American Economic Review, Papers and Proceedings*, 91(2): 67-72.

Blanchflower, David G., and Andrew.J. Oswald (2007). Is Well-being U-Shaped over the Life Cycle?, NBER Working Paper No. 12935.

Browning, Martin, and Thomas F. Crossley (2001) 'The Life-Cycle Model of Consumption and Saving,' *Journal of Economic Perspectives*, 15, 3-22

Brzozowski, Matthew, and Yuqian Lu, (2006), "Home Cooking, Food Consumption and Food Production among the Unemployed and Retired Households," SEDAP Working Paper #151.

Crossley, Thomas and Krishna Pendakur, (2006), "Consumption Inequality", In Green, D. and J. Kesselman, (Eds.), *Dimensions of Inequality in Canada*, UBC Press.

Easterlin, Richard A., (1974), "Does economic growth improve the human lot? Some empirical evidence," in David PA. and Melvin, W.R. (eds.), *Nations and Households in Economic Growth*, Palo Alto, Ca: Stanford University Press, 98-125.

Easterlin, Richard A. (1995), ``Will raising the incomes of all increase the happiness of all?'' *Journal of Economic Behavior and Organization*, 27(1):35-47.

Ferrer-i-Carbonell, Ada, and Paul Frijters, (2004), "How Important is Methodology for the Estimates of the Determinants of Happiness?" *Economic Journal*, 114:641-659

Frey, Bruno S. and Alois Stutzer, (2002), "What Can Economists Learn from Happiness Research?" *Journal of Economic Literature*, 40(2):402-435.

Hurd, Michael D,. and Susann Rohwedder, (2006), "Some Answers to the Retirement-Consumption Puzzle" NBER Working Paper No. 12057

Kahneman, Daniel and Alan Krueger (2006) "Developments in the Measurement of Subjective Well-Being", *Journal of Economic Perspectives*, 20(1): 3-24.

Kapteyn, Arie, Jim Smith and Arthur van Soest, (2006), "An International Comparison of Work Disability", *American Economic Review*, forthcoming.

King, Gary, Christopher Murray, Joshua Salmon and Ajay Tandon (2004), "Enhancing the validity and Cross-Cultural Comparability of Measurement in Survey Research," *American Political Science Review*, 98(1): 191-207.

Robb, A L; Burbidge, J B, "Consumption, Income, and Retirement", *Canadian Journal of Economics*, vol. 22, no. 3, August 1989, pp. 522-42

Scholz, John Karl, Ananth Seshadri, Surachai Khitatrakun, (2006), "Are Americans Saving "Optimally" for Retirement?" *Journal of Political Economy,* 2006, vol. 114, no. 4, pgs 607 - 643.

Smith, Sarah, (2006), "The retirement-consumption puzzle and involuntary early retirement: Evidence from the British Household Panel," *Economic Journal*, 116:C130-C148

Tables

Table 1: Relative Life Satisfaction of Canadian Retirees

Retirement Survey *Is any aspect of life better or worse than expected?*			**General Social Survey** *Compared to the year before you retired, do you now enjoy life more, less, or about the same?*						
	Aged 55 +	**Aged 70+**		**Aged 55 +**			**Aged 70 +**			**Retired within last 5 years**
	(1)	(2)		(3)	(4)	(5)	(6)	(7)	(8)	(9)
	1975	**1975**		**1989**	**1994**	**2002**	**1989**	**1994**	**2002**	**2002**
Worse	23%	21%	**Less**	18%	17%	12%	20%	17%	14%	10%
			Same	39%	34%	46%	40%	36%	53%	41%
Better	25%	32%	**More**	44%	48%	42%	39%	44%	33%	49%

Table 2: Sources of Relative Life Dissatisfaction, Canadian Retirees

Retirement Survey *Aspect of life that is worse than expected*			**General Social Survey** *What is the main reason that you now enjoy life less? Is it....*					
	Aged 55 +	**Aged 70+**	**Aged 55 +**			**Aged 70 +**		
	(1)	(2)	(3)	(4)		(5)	(6)	
	1975	**1975**	**1989**	**1994**	**2002**	**1989**	**1994**	**2002**
Income / Economic	43% [15%]	39% [12%]	7% [1%]	12% [2%]	Not Asked	6% [1%]	6% [1%]	Not Asked
Health	23% [8%]	27% [8%]	58% [10]	46% [8%]		61% [12%]	53% [9%]	
			Conditional % [Unconditional %]					

Notes to Tables 1, 2

The exact wording of the questions in RS75 used here is:

- *Are there any aspects of your life now which are much better than you expected before you retired? Please specify.*
- *Is there anything much worse than expected before you retired? Please specify.*

Table 3: Relative Financial Satisfaction of Canadian Retirees

General Social Survey

Compared to the year before you retired, would you know say that you are better off financially, worse off or about the same?

		Aged 55 +			Aged 70 +		Retired within last 5 years
		(1)	(2)		(3)	(4)	(5)
	1989	1994	2002	1989	1994	2002	2002
Worse		30.5%	26%		23%	21%	31%
Same	Not Asked	50%	56%	Not Asked	51%	57%	57%
Better		19.5%	18%		26%	22%	12%

Table 4: Summary Statistics of Relative Financial and Relative Life Satisfaction

2002 General Social Survey

	(Relative) Life Satisfaction			(Relative) Financial Satisfaction		
Demographics	*Worse*	*Same*	*Better*	*Worse*	*Same*	*Better*
Married/Common Law Male	9.9	46.6	43.5	26.7	57.6	15.8
Married/Common Law Female	9	42.7	48.3	25.9	56	18.1
Widowed/Divorced/Separated Male	18.4	49.3	32.4	29.6	52.8	17.7
Widowed/Divorced/Separated Female	18	47.5	34.4	26.5	52.3	21.2
Single Male	20.3	46.7	33	26.6	54.8	18.7
Single Female	10	44.3	45.7	17.5	55.2	27.3
Age 55-59	9.3	25.6	64.8	28.9	58.5	12.6
Age 60-64	8	42.9	49.1	32.6	54.8	12.6
Age 65-69	12.6	40.7	46.7	31.3	52.9	15.8
Age 70-74	10.9	52	37	25.8	55.3	19
Age 75-79	15.8	49.5	34.8	21.9	57.7	20.4
Age 80+	16.7	58.5	24.8	15.8	57.2	27
Atlantic	10.9	50.2	38.9	19.3	56.9	23.8
Quebec	15.6	39.1	45.3	16.5	60	23.4
Prairies	10.9	50.7	38.3	23.5	56.7	19.8
BC	9.3	50.7	40	31.5	49.6	19
Ontario	12.4	45.6	42	27.8	54.7	17.5
Less than High School	14.7	48.9	36.4	24	59	17
High School Diploma	12	43.5	44.4	28	52	19.8
University graduated	9.5	44.8	45.7	27.6	54.8	17.6
Retirement Reasons						
Mandatory	11.7	47.2	41.1	28.2	55.6	16.2
New technology	12.3	45	42.7	27.2	57.4	15.4
Poor health	19.6	44.1	36.3	35	49.9	15
Unemployed	24.4	49	26.6	36.7	47.3	16
Care Taking	15	52.1	32.9	26	53.1	20.9
Job Downsized	13	45.4	41.6	30.4	54.7	14.9
Business Closure	17.7	49.9	32.5	27.3	58.2	14.5

Continues on next page

Family Reasons	12.3	46.2	41.5	21	50.6	28.4
Voluntary Decision	9.1	45.8	45	22.5	58.5	19
Other						
Home Ownership	9.9	46.1	44	26.6	55.4	18
"Not Home Owner"	20.4	45.5	34	25.8	56.1	18.1
Benefit from Employer	8.9	44.2	46.9	26.3	55.4	18.4
"No Employer Pension"	15.9	47.9	36.2	26.6	55.9	17.5

Notes to Table 4

1. Variables are defined as follows:
   a) *Married*: "Living Common in Law + Married"
   b) *Voluntary Decision*: Was your retirement voluntary, did you want to retire?
   c) *Home Owner:* Is this dwelling owned by a member of this household?
   d) *Benefit from Employer*: Do you receive a pension or retirement pension from any of your former employers?
   e) *Retirement Reasons:* Why did you retire? Was it because...
      Mandatory: Your Employer had a mandatory retirement policy?
      New Tech.: New Technology was introduced?
      Poor Health: Your health required it?
      Unemployed: You were unemployed and could not find a job?
      Care Taking: You needed to take care of a family member?
      Job downsized: Your job was downsized
      Business Closure + Family Reasons are created from inputs of people responded as other reasons
      Business Closure: Of a business closure or a lay off
      Family Reasons: Of a family reasons included re-location

Table 5: Ordered Probit, Relative Life Satisfaction, GSS 2002

	Coefficient	Marginal Effect (More or Same)	Coefficient	Marginal Effect (More or Same)
Age [60-64]	-.274**	-.056	-.302**	-.063
	(-2.72)		(-3.19)	
Age [65-69]	-.373**	-.077	-.403**	-.085
	(-3.72)		(-4.30)	
Age [70-74]	-.524**	-.115	-.539**	-.120
	(-5.26)		(-5.82)	
Age [75-79]	-.638**	-.150	-.637**	-.150
	(-6.16)		(-6.59)	
Age 80+	-.786**	-.199	-.807**	-.206
	(-7.59)		(-8.32)	
High School Grad.	.076	.013	.055	.010
	(1.60)		(1.17)	
College Graduated	.086	.015	.069	.012
	(1.87)		(1.54)	
Atlantic	-.022	-.004	-.015	-.002
	(-0.45)		(-0.31)	
Quebec	.041	.007	.065	.011
	(0.72)		(1.14)	
Prairie	-.027	-.005	-.013	-.002
	(-0.56)		(-0.28)	
BC	.041	.007	.047	.008
	(0.81)		(0.94)	
Married Male	.206	.037	.187	.034
	(1.63)		(1.54)	
Married Female	.275*	.046	.276*	.047
	(2.08)		(2.18)	
Male, Separated,	.044	.008	.023	.004
Divorced or Widowed	(0.33)		(0.18)	
Female, Separated,	.131	.023	.119	.021
Divorced or Widowed	(1.02)		(0.97)	
Never Married Female	.332*	.050	.349*	.052
	(2.21)		(2.37)	

Continues on next page

Home Owner	.173**	.033	.169**	.033
	(3.39)		(3.41)	
Benefit from Employer	.222**	.041	.214**	.040
	(5.68)		(5.51)	
Voluntary Decision	.439**	.093		
	(9.62)			
Mandatory			-.108*	-.021
			(-1.98)	
New Tech.			-.048	-.009
			(-0.52)	
Poor Health			-.314**	-.064
			(-6.97)	
Unemployed			-.425**	-.098
			(-4.61)	
Care Taking			-.195**	-.040
			(-3.08)	
Job Downsized			-.069	-.013
			(-1.14)	
Business Closure			-.247**	-.052
			(-2.99)	
Family Reasons			-.094	-.018
			(-1.13)	
Observations	7060		7112	

Notes to Table 5

1. Variable definitions follow Table 4.
2. z-values in parentheses
3. ** indicates significant at the 1 percent level, * indicates significant at the 5 percent level
4. Reference (omitted) groups for categorical variables are: Age 50-54, Ontario, Less Than high School, Single Male
5. Table 5 is based on responses to the question: "Compared to the year before you retired , do you now enjoy life more, less, or about the same?" The dependent variable is coded 2 for "more", 1 for "about the same", and 0 for "less".

Table 6: Ordered Probit, Relative Financial Satisfaction, GSS 2002

	Coefficient	Marginal Effect (Better or Same)	Coefficient	Marginal Effect (Better or Same)
Age [60-64]	-.050	-.016	-.082	-.027
	(-0.64)		(-1.02)	
Age [65-69]	.015	.005	.002	.000
	(0.21)		(0.03)	
Age [70-74]	.181*	.056	.168*	.052
	(2.43)		(2.21)	
Age [75-79]	.271**	.082	.264**	.080
	(3.57)		(3.39)	
Age 80+	.489**	.138	.477**	.135
	(6.03)		(5.76)	
High School Grad.	-.014	-.004	-.028	-.009
	(-0.32)		(-0.65)	
College Graduated	-.048	-.015	-.056	-.018
	(-1.18)		(-1.38)	
Atlantic	.126*	.039	.129*	.040
	(2.58)		(2.67)	
Quebec	.096*	.030	.109*	.034
	(1.97)		(2.26)	
Prairie	.071	.022	.091	.028
	(1.48)		(1.90)	
BC	-.046	-.015	-.056	-.018
	(-0.87)		(-1.07)	
Married Male	-.077	-.025	-.069	-.022
	(-0.72)		(-0.65)	
Married Female	.049	.015	.030	.009
	(0.45)		(0.27)	
Male, Separated,	-.130	-.043	-.126	-.042
Divorced or Widowed	(-1.13)		(-1.10)	
Female, Separated,	-.026	-.008	-.039	-.012
Divorced or Widowed	(-0.24)		(-0.37)	
Never Married Female	.243	.072	.264	.078
	(1.80)		(1.95)	

Continues on next page

Home Owner	.048	.015	.048	.015
	(1.15)		(1.14)	
Benefit from Employer	.034	.011	.070	.022
	(0.94)		(1.87)	
Voluntary Decision	.332**	.113		
	(7.68)			
Mandatory			-.190*	-.064
			(-3.47)	
New Tech.			-.015	-005
			(-0.17)	
Poor Health			-.258**	-086
			(-6.21)	
Unemployed			-.151	-.051
			(-1.59)	
Care Taking			.013	.004
			(0.22)	
Job Downsized			-.094	-.031
			(-1.60)	
Business Closure			-.048	-.015
			(-0.58)	
Family Reasons			.210*	.063
			(2.48)	
Observations	7067			

Notes to Table 6

1. Variable definitions follow Table 4.
2. z-values in parentheses
3. ** indicates significant at the 1 percent level, * indicates significant at the 5 percent level
4. Reference (omitted) groups for categorical variables are: Age 50-54, Ontario, Less Than high School, Single Male
5. Table 6 is based on responses to the question: "Compared to the year before you retired, would you know say that you are better off financially, worse off or about the same?" The dependent Variable is coded 2 for "better", 1 for "about the same", and 0 for "worse."

Figure

Figure 1: Life Cycle Paths of Income, Consumption and Happiness

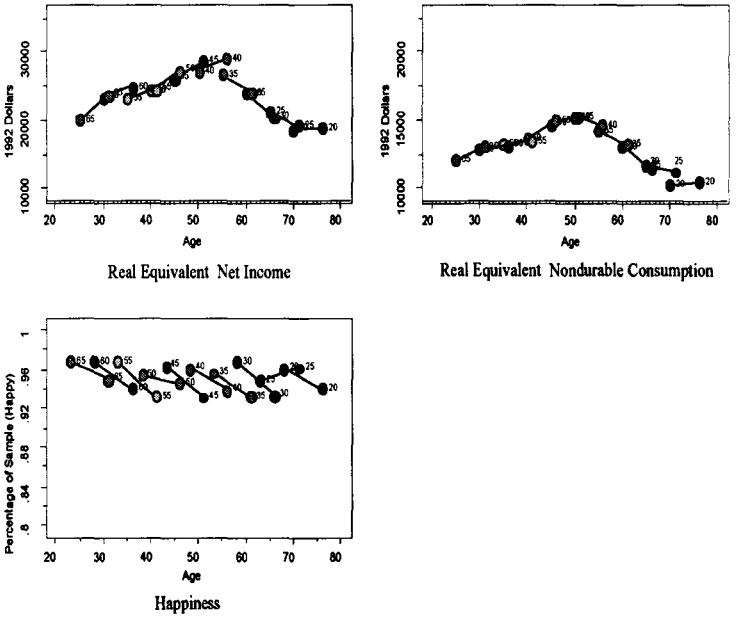

Figure 1: Life Cycle Paths of Income, Consumption and Happiness

Income, Consumption: 1992 Family Expenditure Survey and 1998 Survey of Household Spending
Happiness: 1990 and 1998 General Social Surveys

Appendix

Table A1: Details of the General Social Surveys

	1989 (Cycle 4)	1994 (Cycle 9)	2002 (Cycle 16)
Sampling Frame	Random Digit Dialing Method (RDD).	Random Digit Dialing Method with supplementary sample drawn from Labor Force Survey sampling frame.	Respondents were randomly selected from a list of individuals aged 45 years and over who had responded to the Canadian Community Health Survey (CCHS) in 2001. The CCHS in turned used LFS sampling frame and RDD.
Target Population	Persons *15* years of age and older living in the 10 provinces*.	Persons *15* years of age and older living in the 10 provinces*.	Persons *45* years of age and older living in the 10 provinces †.
Mode	Telephone interview	Telephone interview	Telephone interview
Core Content	Education and Work	Education, Work and Retirement	Social Support and Aging
Sample Size (Response Rate)	9,338 (80.7)	11,876 (81.2)	24,870 (86.3)

Notes to Table A1

* The target population for the GSS was all persons 15 years of age and over in Canada, excluding residents of the Yukon and Northwest Territories and full-time residents of institutions.

† The target population for Cycle 16 of the GSS was all people aged 45 and over in Canada, excluding residents of the Yukon, Northwest Territories, and Nunavut; full-time residents of institutions; residents living on Indian Reserves, Crown lands or in some remote areas; full-time members of the Canadian Armed Forces.

Table A2: Summary statistics, Analysis Sample (Retired Respondents)
2002 General Social Survey

Demographics	*Percentage*
Male, married	40.95
Female, married	23.84
Male, divorced/widowed	8.35
Female, divorced/ widowed	20.8
Male, never married	2.82
Female, never married	3.18
Age 55-59	8.52
Age 60-64	15.49
Age 65-69	22.96
Age 70-74	21.40
Age 75-79	17.23
Age 80+	14.40
Atlantic	8.27
Quebec	22.08
Prairies	14.37
BC	14.75
Ontario	40.53
Less than high school	38.64
High school diploma /+	26.37
University graduate	33.99

Retirement Reasons	*Percentage*
Mandatory	12.60
New technology	3.97
Poor health	23.21
Unemployed	4.87
Care Taking	8.63
Job Downsized	9.92
Business Closure	4.99
Family Reasons	4.89
Voluntary Decision	75.31

Financial Resources	
Home ownership	75.88
Pension benefit from former employer	47.24

N=7940

Figure A1: Inflation and Unemployment in Canada

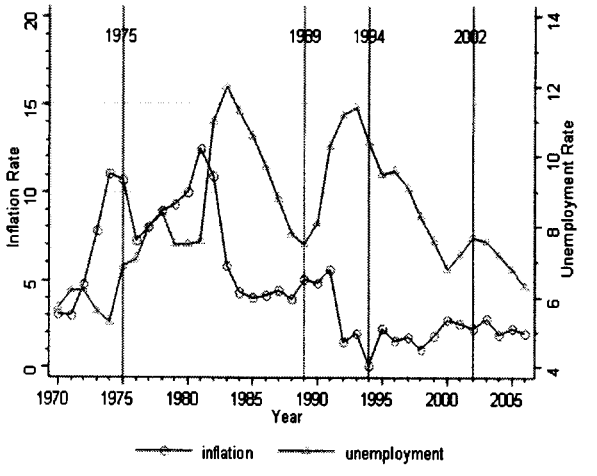

Appendix Figure A1: Inflation and Unemployment in Canada (CANSIM data)

III
Do the Rich Save More in Canada?

1. Introduction

Do the rich save more? This is an important question for a myriad of policy issues, including: Is a switch from income taxation to consumption taxation regressive? What are the consequences of income inequality for economic growth? What is the effect of a tax cut on aggregate demand? What is the incidence of the tax expenditures associated with tax favoured saving accounts?

As Dynan et al. (2004) have recently pointed out, most non-economists would find the proposition that the rich save more to be obvious. Economists are more skeptical, for at least three reasons. 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 current income and saving rates is to be expected, but tells us little about the relationship between saving rates and permanent or lifetime income. Second, is the logic of budget constraints: in the absence of (intended or unintended) bequests, a lifetime budget constraint implies that if a lifetime income group saves more rapidly at some ages, the same group must dissave more rapidly at other ages. Third, most of our standard models assume features (for example, intertemporally 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 lifetime 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 lifetime 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 lifetime earnings and a poor household to have the opposite expectation.) However, direct empirical evidence suggests that desired bequests are small (Hurd, 1987). Finally, 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).

The relationship between lifetime incomes and saving rates remains, then, 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 has returned to the issue of lifetime incomes and saving rates. Bernheim and Scholz (1993) and Hubbard et al (1994) demonstrated that wealth levels are disproportionately high among households with high

lifetime income. Wealth levels, of course, reflect both past rates of (active) saving and past rates of return. Most recently, Dynan et al. (2004) use three different U.S. data sets and several different instruments (including the ones we use in this paper) 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 lifetime) income.

The goal of this paper is to provide some new evidence on this question using Canadian data and methods similar to those employed by Dynan et al. (2004). This analysis is obviously an important input into Canadian policy making, and a useful replication of the Dynan et al. study on data drawn from a different, but similar economic environment. However, additional value flows from particular and unique features of the Canadian data. Saving can be studied using household expenditure surveys (to construct measures of income minus consumption, or active saving) or from panel data on household wealth (possibly with a correction for capital gains to give a measure of active saving). Canadian survey data on household wealth does not have a panel component. However, the Canadian Family Expenditure Survey (FAMEX) has several features that have lead researchers to believe that it can be the basis for a very good measure of active saving. First, in contrast to the U.S. Consumer Expenditure Survey, the FAMEX was particularly designed to capture good quality income information that refers to the same (annual) period as the expenditure information. Second, the FAMEX, in addition to annual income and expenditures, reports net changes in assets and debts over the year, excluding capital gains (by tracking additions to, and withdrawals from, financial assets,

as well as changes in debt). This represents a second measure of active saving for the same households. Partly for these reasons, the FAMEX has formed the basis of a number of studies of saving behavior by both Canadian and U.S. authors[1].

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 lifetime or permanent income to employ these data. Our empirical analysis suggests that the estimated relationship between saving rates and lifetime incomes is sensitive to the instrument used to proxy lifetime income. Nevertheless, our preferred estimates indicate that, except for the poorest quintile of households (who simply do not save), saving rates do not differ substantially across lifetime 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

The FAMEX is a full household expenditure survey (collecting information on all categories of expenditure). Unlike most national expenditure surveys, the FAMEX does not have a diary component. Instead, face-to-face interviews are 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 are asked to consult bills and receipts and if

[1] See for example: Burbidge and Davies, 1994; Carroll et al., 1994; Davies and Burbidge, 1994; Engelhardt, 1996; Burbidge et al., 1998; Lin, 2000; Veall and Fretz, 2000; Milligan, 2002.

necessary, multiple visits are made to a household. The FAMEX is therefore an unusual kind of recall survey in which a considerable effort is made to ensure the quality of the data.

Our analysis is based on public use files from the 1996 survey. 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 lifetime incomes, the appropriate saving concept is active saving (or the "true" saving intention.) An important feature of the FAMEX is that it contains two measures of household active saving. The first is simply after-tax current income minus total expenditure. (This measure is also used in Dynan et al (2004) for CEX and PSID; in the case of the latter total expenditure must be imputed). As noted above, this measure may be of higher quality in the FAMEX than CEX because of the nature of the data collection exercise.

The second measure of active saving in the FAMEX is net changes in assets and debts excluding capital gains. This measure is unique (not directly available in any of the American Surveys used by Dynan et al.) 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-favoured individual retirement account); net purchases less sales of financial assets; sales of personal property.

In the conduct of the survey these two measures are partially reconciled in that household in which the two measure show an excessive disparity are asked to review their reports of incomes and expenditures. Consequently, measurement errors are unlikely to be independent across the two measures. At the same time, the second measure appears to contain additional information. 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 the FAMEX is superior to those in the CEX (where the income data is not ideally suited to this purpose) and PSID (in which total expenditure must be imputed), and the FAMEX contains a second measure of annual active saving that can be exploited in a number of ways (which we outline below)[2].

There are 10085 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. The first restriction is 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 regarding educational choice and dissaving in retirement. We also exclude households that reported less than $1000 of income and households that did not report their education level. Finally, we delete multiple family units (more than 1

[2] Against this, unlike Dynan et al., we do not have panel data on wealth. 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. Change in stock of wealth can be only be calculated with 5 year intervals (PSID wealth supplements are 5 year apart panel surveys conducted in 1984, 1989, 1994, 1999 and 2003). In both cases, changes in wealth must be purged of capital gains to construct a measure of active savings.

family living in the same dwelling), which are a small fraction of the sample[3].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 (ΔA) and then we divide by current income to give the saving rate ($\frac{S}{Y}$). 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 civil war pensions and allowances) and income from other sources (alimony, RRSP annuities received and RRIF withdrawals, child support). 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, contributions and the interest portion of mortgage payments as consumption. The portion of mortgage payments that is principal repayment is treated as saving. Note that individual contributions to tax sheltered savings plans (RRSPs) are counted as saving, while contributions to public and employer sponsored pensions are not (neither employee nor employer contributions to these plans are counted in net income.)

[3] 3% of the full sample and 1% of the couples only sample.

Descriptive statistics for these measures are given in Table 1. The first thing to note is that the (Y-C) measure of saving suggests higher levels and rates of saving. This is consistent with the under-reporting of consumption or the under-reporting of contributions to financial assets and/or retirement of debt. The second thing that we note is 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, in this sense, on average richer than the full sample. Thus the lowest quintile of the full sample has lower incomes than the lowest quintile of the sample of couples[4].

3. Empirical Methodology

We wish to estimate the relationship between saving rates and lifetime income:

$$\frac{S}{Y} = f(Y^P) + X\beta + e \qquad (1)$$

where Y^P is lifetime 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. To allow for nonlinearities in the relationship between lifetime incomes and saving rates, we parameterize $f()$ by a set of five dummies capturing the age-conditional quintile of lifetime income to which a each household belongs (X does not include a constant).

[4] Following Dynan et al., we have not made any adjustments for household size or compositions to income or consumption (for example, converting to per capita amounts or dividing by an equivalence scale).

The key empirical problem we face is that we do not observe lifetime income (Y^P). Moreover, for our purposes current income, (Y) is a poor proxy for lifetime income because the smoothing of transitory income fluctuations will generate a positive relationship between saving and current income and even when there is no relationship with lifetime incomes. Our solution, which follows Dynan et al. (2004) is a two stage estimation procedure. In the first stage we construct lifetime income proxies by regressing current income on instruments (Z) and age group dummies:

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

Predicted values from this regression are then used as our proxy for lifetime income.

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

We then assign households to age-conditional lifetime income quintiles, and construct the quintile dummies that were described above. In the second stage we estimate Equation (1) by quantile regression. Since lifetime income is estimated in the first stage we bootstrap the standard errors[5].

One way that we can 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 temporarily separated - they refer to the same year.) In principle this could increase the precision of our estimates, and so we report estimates based on pooling the data below, along with estimates based on each measure separately. When we pool the data we allow for a common mean shift between the two measures, and we take care in our bootstrapping to resample households (pairs of

[5] Bootstrap standard errors are based on 999 replications.

observations) in order to preserve the correlation structure in the data (as in panel data bootstrapping). In practice, this does not lead to much increase in precision. However, a second way in which we can exploit the second measure of active saving in the data is to strengthen our strategy for proxying lifetime income, as is discussed below.

The key to our empirical strategy is obviously the instruments for lifetime income. These must be (i) strongly correlated with life-time 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 instruments for lifetime income that are also employed by Dynan et al.: education and nondurable consumption (or components of nondurable consumption)[6]. There is not much doubt that both these instruments are strongly correlated with lifetime income. However, the second condition may be violated for reasons specific to each instrument.

Although it is highly correlated with lifetime 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' discount rates; impatience is also associated with lower lifetime saving[7]. If education is related to preference heterogeneity that is important for saving behavior, then it is an invalid instrument (because it is correlated with the error term in Equation (1)). The likely consequence of this is an upward bias in the estimated relationship between

[6] In the parts of their analysis that are based on panel data, Dynan et al. have two additional instruments that are not available to us: lagged and future earnings.

[7] Similarly, it might be argued that educational choices and savings are both driven in part by heterogeneity in risk aversion.

lifetime income and saving rates (the patient accumulate more education and save more). The information on education in the FAMEX is categorical (less than 9 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 lifetime 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 an instrument are for the sample of couples only.

We define nondurable consumption as total consumption minus spending on shelter, vehicles and household furnishings[8]. If we use nondurable consumption (C^n) to proxy lifetime 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.

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

Lifetime income is positively correlated with nondurable consumption, and consumption enters the saving rate negatively. If the true relationship between saving rates and lifetime income is positive, then measurement error common to nondurable and total consumption will impart a negative bias to our estimates, biasing them towards zero. The same

[8] Thus it contains spending on food, household operations, cloth, health care, personal care, tobacco & alcoholic beverages, reading materials and miscellaneous expenses, plus transportation and recreation minus purchases of cars and recreational vehicles.

problem arises if consumption has a "transitory component" (for example, 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 the FAMEX, we can replace (Y-C) by our second measure of saving, (ΔA) in our estimating equation. To the extent that measurement errors in (ΔA) are not perfectly correlated with measurement errors in consumption, this should reduce the bias. Second, rather than use (all) nondurable consumption as our instrument, we can use well-measured components of nondurable consumption. Food is one possibility. Other possibilities are items that are regularly billed (as bills can be consulted during the survey) but are lifetime income elastic. Discretionary telecommunications expenses (phone bills) are one example. Again this strategy reduces bias by minimizing the potential correlation between measurement errors in our saving rate and our instrument.

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

4. Results

Recall that in all our median regressions we suppress the constant and include dummies for all five lifetime 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 lifetime income quantile whose head is between 40 and 49 years old.

We begin our analysis by documenting the estimated relationship between saving rates and current incomes. Table 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 from pooling the two measures (but allowing for an intercept shift.) 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 three lines (corresponding to estimates based on (Y-C)/Y, (ΔA)/Y and pooling the two) and five points on each line (corresponding to the five income quintiles).

The results confirm that savings rates are strongly increasing in current income. For example, focusing on the full sample and the (ΔA)/Y saving measure, median saving rates for 40 to 49 year old households range from 0 percent in the lowest income quintile to 16 percent in the highest quintile; the corresponding numbers are -6.3 percent to 27.4 % 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%); of course, current incomes are more disperse in the U.S. data. The

stars on quantile 2 through 5 coefficients in Table 2 indicate that each coefficient is statistically significantly different from the coefficient for the quantile below it (at the 5% level).

We now turn to the relationship between saving rates and lifetime incomes, which is our primary interest. Table 3 reports estimated saving rates by life-time income quintiles (from median regressions.) All the estimates in this Table are based on the couples sample. The first three columns of Table 3 (on the left) give results using the education of the head and spouse as instruments for lifetime income. The three columns report estimates based on (Y-C)/Y, (ΔA)/Y and pooling the two (moving from left to right). These results are summarized in Figure 2. The format of Figure 2 (and subsequent Figures) is the same as Figure 1 except that each point represents a quintile of lifetime income. The last three columns of Table 3 (on the right) give results using nondurable consumption as the instrument (with saving measures based on (Y-C)/Y, (ΔA)/Y and pooling the two.) These results are summarized in the right panel of Figure 3. The left panel of Figure 3 also summarizes results using nondurable consumption as the instrument, but for the full sample. (The estimates underlying the (ΔA)/Y line in this graph are given in the fourth column of Table 4; full results are available from the authors).

The first aspect of these results to note is that the choice of saving measure ((Y-C)/Y or (ΔA)/Y) makes little difference. Estimates based on (Y-C)/Y give higher saving rates in each quintile than those based on (ΔA)/Y, and estimates based on pooling the two

lie in between. However, the pattern across quintiles is quite similar regardless of choice of measure. In what follows, we focus on the estimates based on (ΔA)/Y.

Using education as the instrument for lifetime income results in a strong positive relationship between saving rates and lifetime incomes. The estimated median saving rate for a 40-49 year old household rises monotonically from 5.6 percent in the bottom quintile of lifetime incomes to 13.6 percent in the top quintile (2nd column of Table 3 and Figure 2). While no quantile coefficient is statistically different from one just below it, the coefficient on the top quantile dummy is strongly statistically different from coefficient on the bottom quantile dummy.

In contrast, when we use nondurable consumption as an instrument, the estimated relationship between saving rates and lifetime incomes is essentially flat. The estimated median saving rate for a 40-49 year old household is 6.4 percent in the bottom quintile of lifetime incomes and 7.2 percent in the top quintile. It actually peaks (at 9.9 percent) in the 2nd quintile.

Figure 3 illustrates an important distinction between the sample of couples (with and without children) and the full sample. For the couple sample, the median estimated lifetime income in the bottom quintile is 33,288 1996 Canadian dollars. For the full sample, the corresponding number is 22,367. This is because the 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 lifetime income distribution. Using the (ΔA)/Y measure, the estimated median saving rate for a 40-49

year old household in the bottom quintile of lifetime incomes in the full sample is 0 (Figure 3 and Column 4 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 lifetime incomes when we use nondurable consumption as an instrument for lifetime incomes is consistent with the US 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 US study, this seems a 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.) To push this further, we replace nondurable consumption as our instrument with components of nondurable 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 4 (for the full sample) and Figure 5 (couples). All of these estimates use $(\Delta A)/Y$ as the measure of the saving rate. In Figure 4 we also include the estimated relationship between saving rates and current incomes for comparison. In Figure 5 we include, for comparison, the estimated relationship between saving rates and current incomes, and the estimated relationship between saving rates and lifetime incomes when education is used as the instrument.

The main message of these results is that the estimated relationship between saving rates and lifetime incomes is not sensitive to whether we use nondurable

consumption as an instrument or a component of nondurable consumption[9]. 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 consumption measure as an instrument for lifetime income results in an estimated relationship between saving rates and lifetime incomes that is much flatter than the estimated relationship between saving rates and current incomes. In contrast, when we use education as an instrument for lifetime incomes, the estimated relationship between saving rates and lifetime incomes that is as steep as the estimated relationship between saving rates and current incomes. This can be seen clearly in Figure 5.

5. Conclusion

To summarize, when we use education as an instrument for lifetime income, we find a strong positive relationship between saving rates and lifetime income. Indeed, these results suggest that relationship between saving rates and lifetime income is as steep as the relationship between saving rates and current income. This would be surprising, as it is likely that part of correlation between saving rates and current incomes reflects the smoothing of transitory income shocks.

In contrast, when we use consumption as an instrument for lifetime income, we find that above the bottom lifetime income quintile, saving rates are fairly flat. A concern with these results is that measurement error in consumption imparts a negative bias to the

[9] In fact, we tried a number of components of nondurable consumption beyond those reported here and they also led to similar results.

estimated relationship. However, when we take steps to mitigate this bias (constructing saving from net changes in assets in debts rather than income minus consumption; and using well measured components of nondurable consumption as instruments) we observe very little change in the estimated relationship. This 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 instrument, probably because it is correlated with unobserved tastes for saving. The best guide to the relationship between saving rates and lifetimes incomes are the estimates which use consumption as an instrument for lifetime income. We therefore conclude that the rich do not save more - at least compared to those in the middle of the lifetime income distribution. Saving rates are very flat above the bottom quantile of lifetime 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 nondurable consumption as an instrument. We are able to do so because of the quality and unique features of the FAMEX data.

Our bottom line then 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 for their failure to replicate the saving behavior of the poorest quintile. 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 dissave. This provides a useful guide for future research priorities.

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 of the reduction rate in the Guaranteed Income Supplement (GIS) and income taxes mean that many seniors of modest means will face tax rates of 100 percent or more on income from RRSPs (tax-favoured retirement saving accounts). Thus, these households may have very little retirement saving motive. Alternatively, do the very low savings rates of the poor reflect something about preferences (such as a "consumption floor") or about behaviour (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 behaviour of lower income households remains important.

References

Bernheim , B. D. and J. K. Scholz, (1993), "Private Saving and Public Policy" in J.M. Poterba, editor, "Tax Policy and the Economy," (vol. 7), Cambridge, Mass.: MIT Press.

Burbidge, J. B. and J. B. Davies, (1994), "Government Incentives and Household Saving in Canada", in J.M. Poterba, editor, Public Policies and Household Saving, University of Chicago Press for the NBER, 19-56.

Burbidge, J. B., D. Fretz, and M. R. Veall, (1998), "Canadian and American Saving Rates and the Role of RRSPs," *Canadian Public Policy*, 24(2), 259 - 263.

Carroll, C.D., B.-K. Rhee and C. Rhee, (1994), "Are There Cultural Effects on Saving? Some Cross-Sectional Evidence", *Quarterly Journal of Economics*, 109(3), 685-99.

Davies, J. B., and J. B. Burbidge, (1994), "Household Data on Saving Behavior in Canada" in J.M. Poterba, editor, International Comparisons of Household Saving, University of Chicago Press for the NBER, 11-56.

Dynan K. E., J. Skinner, and S. P. Zeldes, (2004), "Do the Rich Save More?" *Journal of Political Economy*, 112, 397-444.

Engelhardt, G.V., (1996), "Tax Subsidies and Household Saving: Evidence from Canada," *Quarterly Journal of Economics*, 111(4), 1237-1268.

Friedman, M, (1957), A Theory of Consumption Function, Princeton University Press.

Hubbard, R. G., J. Skinner and S. P. Zeldes, (1994), "Precautionary Saving and Social Insurance," *Journal of Political Economy*, 103, 360-99.

Hurd, M. D., (1987), "Savings of the Elderly and Desired Bequests" *American Economic Review*, 77(3), 298-312.

Kennickell, A. B. and L. Woodburn, (1997), "Weighting Design for the 1983-89 SCF Panel," Washington, DC: Federal Reserve Board of Governors.

Lin, X., (2000), "Saving before and after Retirement: A study of Canadian Couples." in F. Denton, D. Fretz and B.G. Spencer, Independence and Economic Security in Old Age, University of British Columbia Press, 215-254.

Mayer, T., (1966), "The Propensity to Consume Permanent Income," *American Economic Review*, 56, 1158-77

Milligan, K., (2002), "Tax-preferred Savings Accounts and Marginal Tax Rates: Evidence on RRSP Participation," *Canadian Journal of Economics*, 35, 436-456.

Schillington, R., (2003), "New Poverty Traps: Means-testing and Modest-Income Seniors," C.D. Howe Institute Backgrounder No. 65, C.D. Howe Institute.

Veall, M.R. and D. Fretz, (2002), "The Effect of the Tax-Transfer System on Retirement Savings," in F. Denton, D. Fretz and B.G. Spencer, Independence and Economic Security in Old Age, University of British Columbia Press, 346-372.

Tables

Table 1: Descriptive Statistics (1996 FAMEX)

Variable	Median	Mean	Std Dev.	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	-2,1391	196,428
Gifts	628	1,610	4,431	0	176,280
Net Income	41,000	44,683	25,194	1,000	251,000
Change in Assets	2,157	3,911	15,733	-222,064	193,000
Savings	4,127	6,095	16,147	-227,447	163,996
Couples – 4205 Households					
Gross Income	61,000	66,187	35024	2000	292,400
Total Consumption	39,134	42,280	18,714	4,717	196,428
Gifts	762	1,627	3,126	0	85,860
Net Income	48,000	51,970	24,596	1,000	251,000
Change in Assets	4,000	5,089	16,566	-132,500	180,335
Savings	6,594	8,062	17,106	-121,646	163,996

Note to Table 1:

1. 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.

Table 2: Median Regression of Saving Rates on Age and Current Income Quintile Dummies

Sample	Couples (4205 households)			Full (6062 households)		
Saving Measure	(Y-C)/Y	(ΔA)/Y	Pooled	(Y-C)/Y	(ΔA)/Y	Pooled
Quintile 1	-3.33	-.15	-3.25	-6.30	0	-3.78
	(1.25)	(.51)	(.93)	(.91)	(.22)	(.74)
Quintile 2	11.23*	5.51*	5.64*	7.84*	2.94*	3.19*
	(1.07)	(.91)	(.87)	(.94)	(.62)	(.74)
Quintile 3	14.54*	9.15	9.17*	13.66*	7.51*	8.43*
	(1.07)	(.77)	(.78)	(.90)	(.71)	(.71)
Quintile 4	20.40*	12.67*	13.65*	18.84*	10.51*	12.32*
	(.86)	(.86)	(.69)	(.87)	(.62)	(.74)
Quintile 5	27.86*	17.10*	19.17*	27.40*	16.36*	19.22*
	(.96)	(.86)	(.90)	(.88)	(.62)	(.68)
Age 30 -39	-.65	.15	-.24	-1.87	0	-.33
	(.97)	(.56)	(.74)	(.87)	(.24)	(.64)
Age 50 –59	2.83	.21	1.70	1.28	0	.86
	(1.12)	(.74)	(.91)	(.96)	(.26)	(.77)
(Y-C)/Y (dummy)	-	-	5.43	-	-	3.69
			(.37)			(.33)

Notes to Table 2:

1. Standard errors based on 999 bootstrap replications
2. For the pooled estimates, the sample size is doubled and the (panel) bootstrap involves resampling pairs of observations.
3. * denotes that the coefficient on this quintile is statistically different than the coefficient on the preceding quintile at the 5 % level.

Table 3: Median Regression of Saving Rates on Age and Lifetime Income Quintile Dummies (Couples)

Instrument(s) for Lifetime Income	Education of Head and Spouse			Nondurable Consumption		
First Stage R^2		0.15			0.49	
Saving Measure	(Y-C)/Y	(ΔA)/Y	Pooled	(Y-C)/Y	(ΔA)/Y	Pooled
Quintile 1	11.04	5.64	4.48	13.60	6.48	6.95
	(.95)	(.94)	(.83)	(1.20)	(.92)	(.94)
Quintile 2	13.23	6.67	6.57	17.90*	9.94*	10.58
	(1.26)	(.98)	(.96)	(1.31)	(.80)	(.92)
Quintile 3	15.38	8.78	8.60	15.38	9.16	8.77
	(1.13)	(1.00)	(.89)	(1.26)	(.91)	(.90)
Quintile 4	19.73*	11.26	12.03	15.54	8.79	9.49
	(1.23)	(.96)	(.98)	(1.15)	(.92)	(.84)
Quintile 5	22.29	13.65	14.14	11.43	7.19	6.10
	(1.39)	(.91)	(.94)	(1.23)	(.95)	(.95)
Dummy	-	-	6.63	-	-	6.26
			(.37)			(.38)
Age 30 -39	-1.39	.08	-.41	-1.55	.18	-.62
	(.96)	(.86)	(.80)	(1.13)	(.76)	(.82)
Age 50 –59	2.14	.12	1.41	2.54	1.04	1.81
	(1.11)	(.95)	(.93)	(1.19)	(.92)	(.89)

Notes to Table 3:

1. Sample size: 4205 households
2. Standard errors based on 999 bootstrap replications
3. For the pooled estimates, the sample size is doubled and the (panel) bootstrap involves re-sampling pairs of observations.
4. * denotes that the coefficient on this quintile is statistically different than the coefficient on the preceding quintile at the 5% level

Table 4: Median Regression of Saving Rates on Age and Lifetime Income Quintile Dummies ((ΔA)/Y Saving Measure)

	Couples (4205 households)			Full (6062 households)		
Instrument(s) for Lifetime Income	Nondurable Consumption	Food	Food and Telephone	Nondurable Consumption	Food	Food and Telephone
First Stage R^2	0.49	0.25	0.27	0.59	0.38	0.39
Quintile 1	6.48	7.55	7.35	0	0	0
	(.92)	(.84)	(.81)	(.10)	(.21)	(.19)
Quintile 2	9.94*	9.37*	10.11*	6.69*	5.69*	6.41*
	(.80)	(1.01)	(.97)	(.83)	(.78)	(.75)
Quintile 3	9.16	9.63	9.71	9.30	9.22	8.78
	(.91)	(.78)	(.75)	(.66)	(.72)	(.80)
Quintile 4	8.79	7.54	7.60	8.54	8.78	8.92
	(.92)	(1.02)	(.92)	(.62)	(.63)	(.59)
Quintile 5	7.19	6.85	6.45	7.46	6.89	6.76
	(.95)	(.87)	(.96)	(.61)	(.70)	(.74)
Age 30 -39	.18	.59	.62	0	0	0
	(.76)	(.80)	(.82)	(.22)	(.36)	(.34)
Age 50 –59	1.04	1.25	1.31	0	.02	0
	(.92)	(.92)	(.85)	(.29)	(.47)	(.40)

Notes to Table 4:

1. Standard errors based on 999 bootstrap replications
2. * denotes that the coefficient on this quintile is statistically different than the coefficient on the preceding quintile at the 5% level.

Figures

Figure 1: Current Income Quintiles and Median Saving Rates

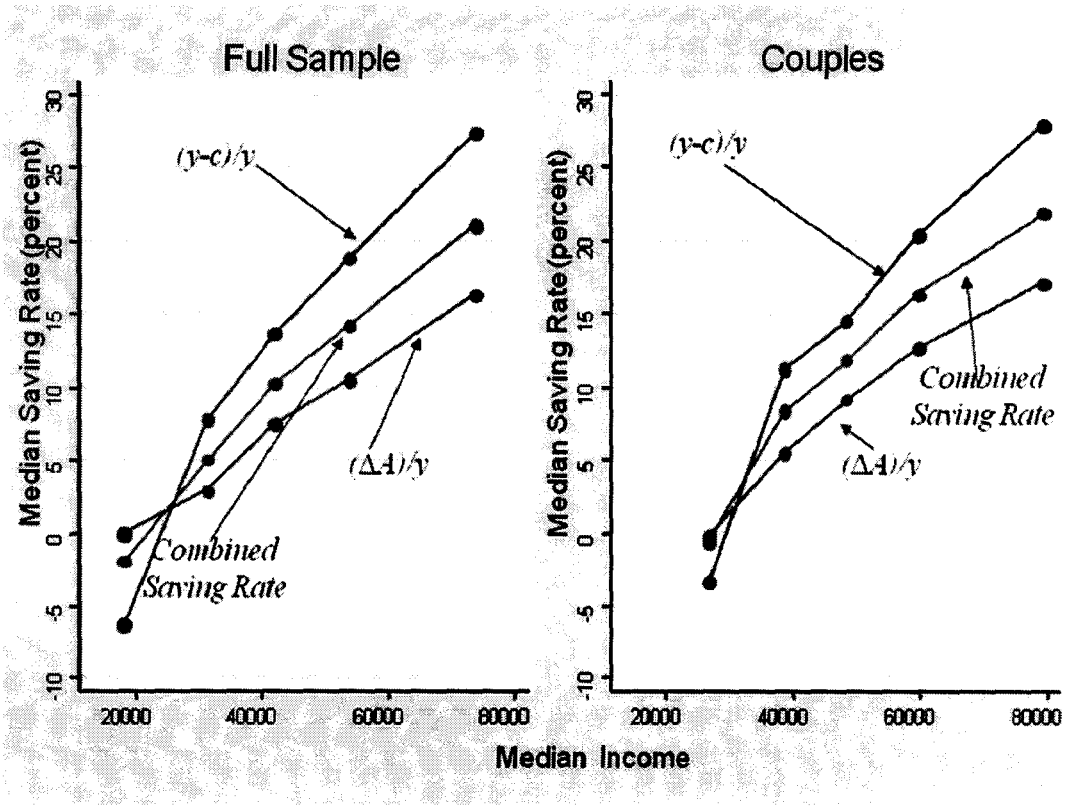

Figure 2: Median Savings Rates and Lifetime Income Quantiles (Couples, Education Instrument)

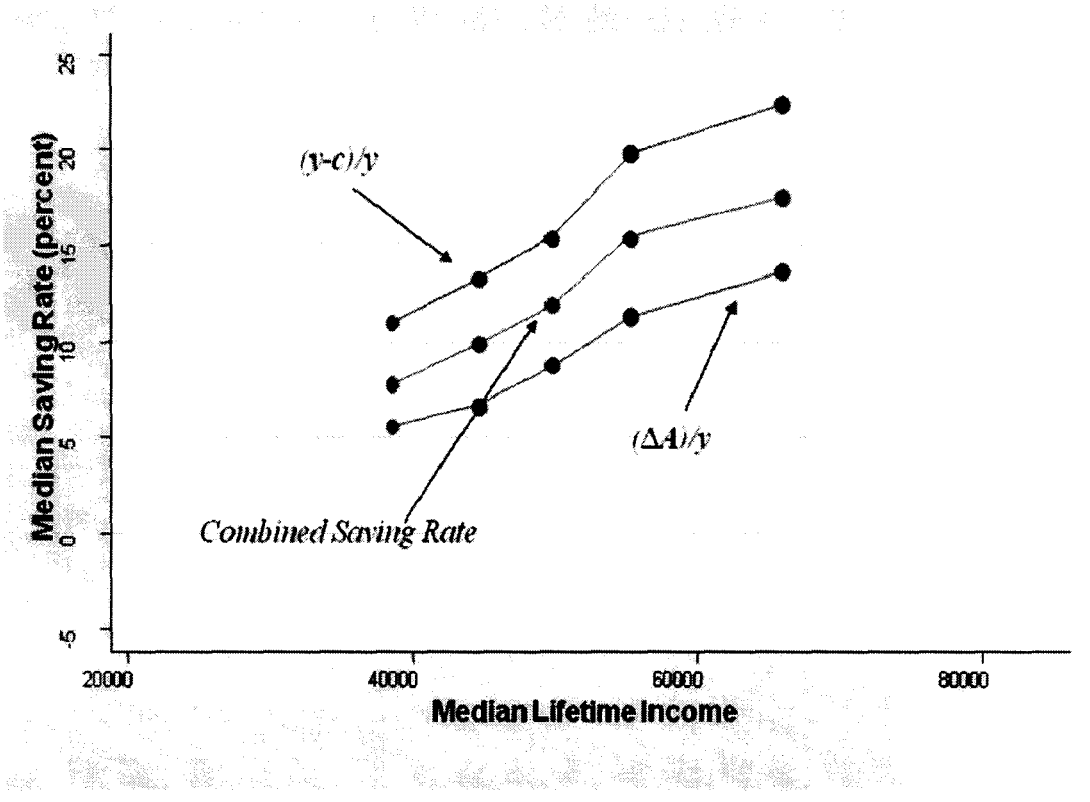

Figure 3: Median Savings Rates and Lifetime Income Quantiles (Nondurable Consumption Instrument)

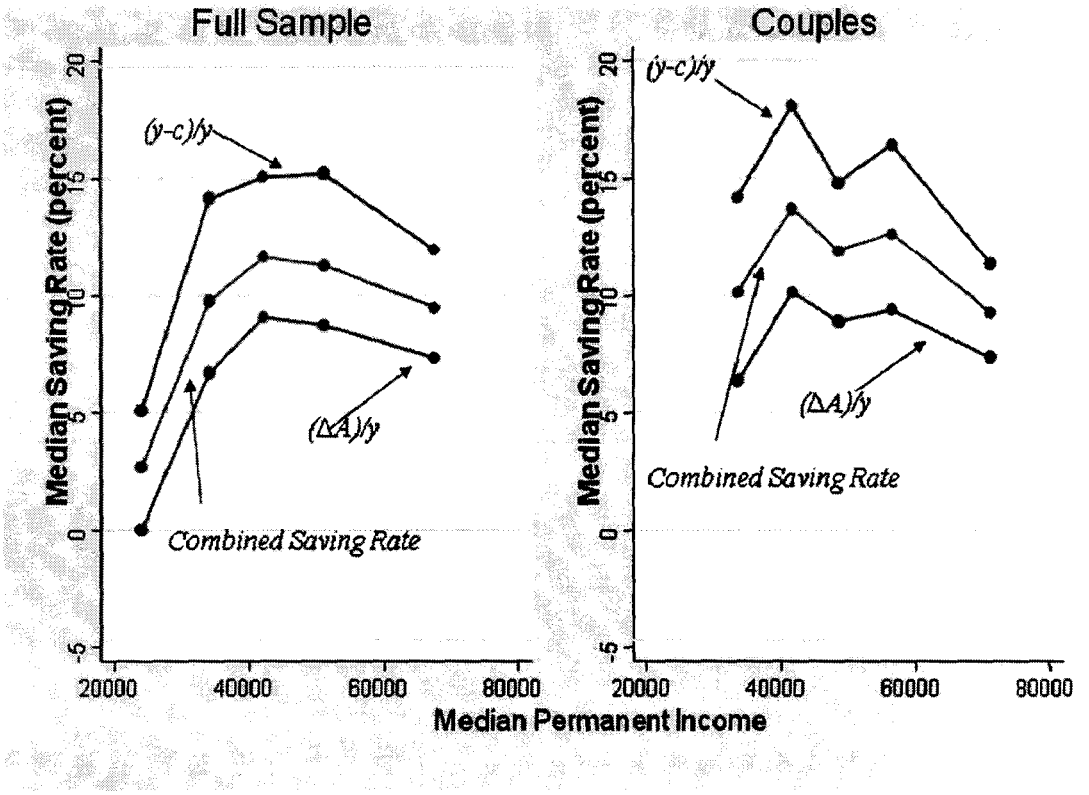

Figure 4: Median Savings Rates and Lifetime Income Quantiles (Full Sample, (ΔA)/y Saving Measure, Alternative Instruments)

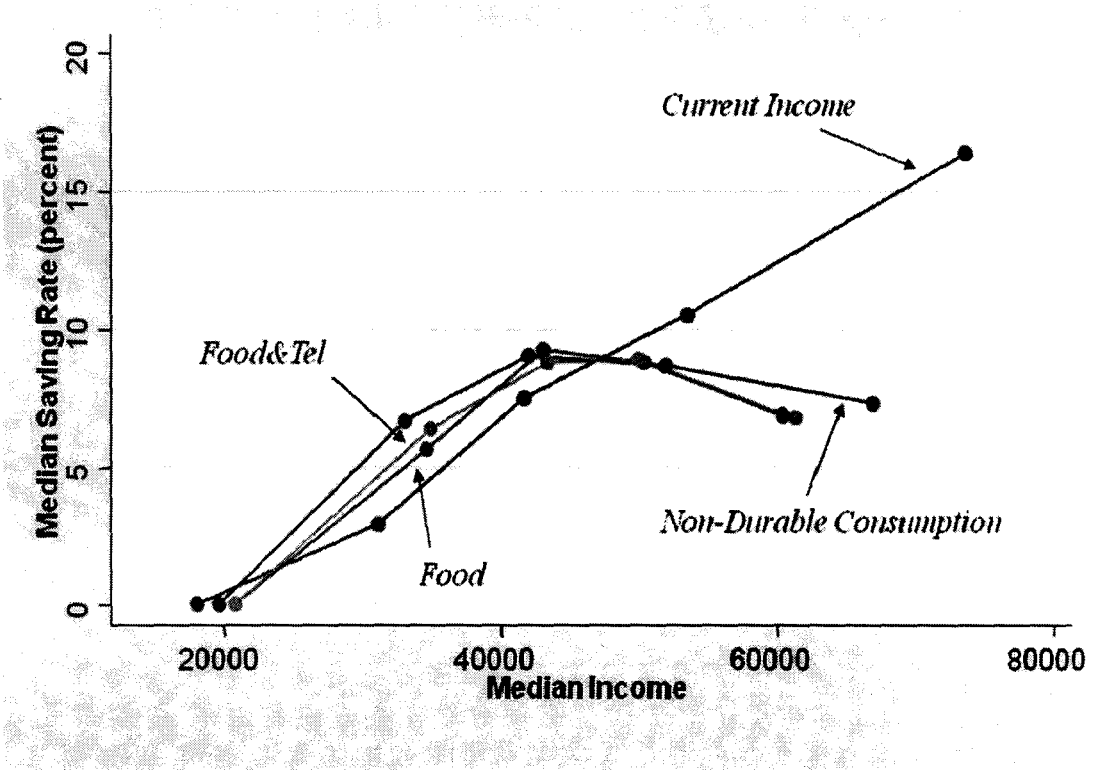

Figure 5: Median Savings Rates and Lifetime Income Quantiles (Couples, (ΔA)/y Saving Measure, Alternative Instruments)

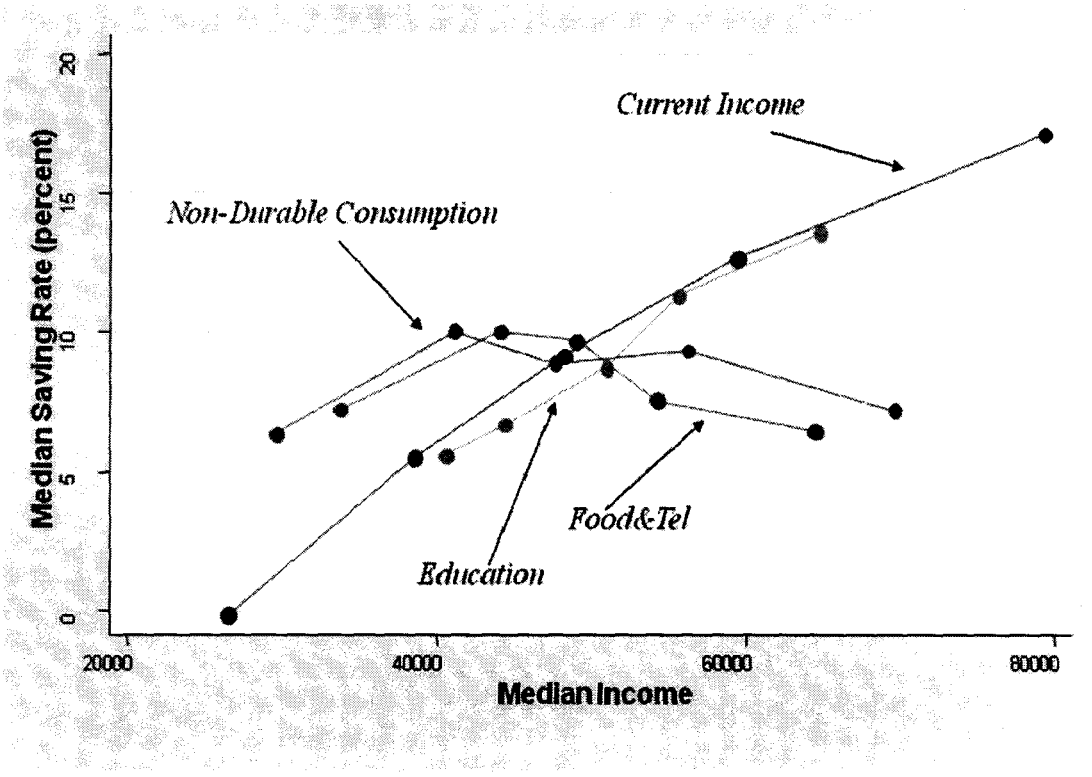

IV

New Evidence on Taxes and Portfolio Choice

1. Introduction

The effect of taxation on household portfolio choice has long been an important question facing researchers and policy makers. Theoretical models predict that under a differential taxation system, households' portfolio allocation decisions are based not only on the risk-return relationship of assets but also on their tax characteristics. Given risk and return characteristics, households should minimize tax burden by tilting their portfolios toward less heavily taxed assets. Moreover, less taxed households face a smaller incentive to invest in tax favored assets. For these theoretical predictions to be a useful guide to tax policy, empirical evidence is required that confirms the qualitative predictions of theory, and which provides quantitative estimates of the magnitudes of relevant effects.

The literature on tax avoidance (Slemrod, 2002) suggests that tax elasticities may be much larger for aspects of financial arrangements (for example, the timing of income) than for real variables (such as labor supply or savings levels). Portfolio choice may be in the former category. Nevertheless, there have been relatively few empirical studies of the effect of taxation on portfolio allocation. The existing literature is well surveyed by Poterba (2001). The key challenge in this literature is to find a substantial and plausibly exogenous source of variation in tax rates. The contribution of this paper is to derive new

estimates of tax effects on portfolio allocation using a novel source of variation in tax rates.

Cross-sectional differences in the marginal tax rates (MTRs) generate variation across households in the relative after-tax returns of different assets classes. Examples of this kind of study include Feldstein (1976), Hubbard (1985) and King and Leape (1998). All these studies find a strong effect of taxes on asset allocation. A problem with this approach however, is that household MTRs are highly correlated with income, and so it is difficult to disentangle a pure tax effect from income or wealth effects on portfolio allocation.

Another strategy, pursued first by Scholz (1994), then Samwick (2000) and Poterba and Samwick (2003) is to study changes in portfolio allocation around tax reforms (using a 'diff-in-diff' approach). The great advantage of this approach is that the tax reform generates variation in MTRs *within* income groups. Moreover, if a suitable control group can be identified, it is possible to control for general time effects, under a 'common trends' assumption. The common trends assumption may also be a weakness of this approach, in that it may be difficult to identify, with any confidence, control groups that were unaffected by the tax reform, but experience time effects that are similar to those that were affected by the reform. Moreover, the investigator faces a difficult tradeoff in deciding the interval over which the data should be 'differenced'. A strategy that compares periods just before and after the reform risks missing delayed adjustments to changes in taxation (households may not rebalance portfolios instantaneously, because of transaction costs or other sources of inertia.) On the other hand, a strategy of

comparing data from long before a tax reform with data well after the tax reform rests more heavily on the common trends assumption and so suffers from a greater risk of confounding a tax effect with other time effects.

It is well known that the last few decades witnessed significant trends in household portfolio allocation. Prior to 1980s most households' financial wealth was held in simple forms (mostly in liquid and safe assets) in most industrialized countries. This observation has changed considerably since 1990s; now a large proportion of households in these countries hold significantly more sophisticated portfolios. Financial liberalization, declining information costs, attraction of employer-sponsored retirement accounts (such as 401Ks in the U.S) and introduction of tax advantaged investment tools (such as registered education saving accounts in Canada) are among the explanations offered for this trend[1]. The large size of time effects in portfolio allocations makes the common trends assumption particularly worrying.

In light of all this, it would be useful to have additional estimates of tax effects on portfolio allocation, based on alterative sources of variation in tax rates. Our proposal is to exploit variation in MTRs across households with the same *total* earnings, which arise in progressive income tax systems with individual taxation. In jurisdictions with individual taxation, such as Canada, two households with the same total earnings, but divided differently between the principal and secondary earner, may face a different MTR on the first dollar of household capital income. In particular, households in which most of the labor income is earned by one individual will face a lower MTR on the first

[1] See Bertaut (1998) and Guiso, Haliassos and Japelli (2002) for exhaustive surveys of household portfolio facts and trends.

dollar of capital income than a household with fairly equal income shares. This is because the former household can attribute capital income to the household member with lower labor earnings. This advantage does not exist in systems of joint taxation, as in the United States. Thus in Canada (and other countries with individual taxation) it is possible to study the effect of MTR on portfolio allocation while holding constant household income and wealth. We use the 1999 Canadian Survey of Financial Security to implement this strategy. Of course, there are a number of important challenges to the validity of this empirical strategy. We address these empirically.

First, this strategy rests on the assumption that in systems with individual taxation, households shift financial assets to the secondary earner (or lower-earning spouse) in order to minimize the taxation of capital income. Stephens and Ward-Batts (2004) report evidence in support of this proposition from a study of the 1990 change from joint to individual taxation in the U.K. Below we provide further evidence supporting this assumption, by studying the distribution of capital income within households before and after the Canadian tax reform of 1988. For many households, that tax reform had the effect of making the Canadian tax system less joint.

Second, households in which labor earnings are fairly equally contributed may differ in important ways (in addition to effective MTR) from households that have similar total labor income and wealth, but greater inequality in labor income shares. For example, households with two labor incomes of fifty thousand dollars annually may have different preferences (including risk tolerance) than a household with a single labor income of one hundred thousand dollars per year. Alternatively, intra-household decision making may

proceed quite differently in these two (for example, bargaining power may be more evenly distributed in the first household than in the second).

To address these concerns, we implement a 'placebo test'. In particular, we study the relationship between labor income shares and portfolio allocation (holding wealth and total income constant) in two U.S. data sets: the 1998 Survey of Consumer of Finances (SCF) and the 1999 wealth module of the Panel Study of Income Dynamics (PSID). Because the U.S. has joint taxation, the first dollar MTR on capital income is unaffected by the distribution of labor income within the household. Thus a correlation between labor income shares and portfolio allocation in these data would suggest important heterogeneity in preferences or household bargaining, while the absence of such a correlation would support our empirical strategy for identifying tax effects.[2]

Veall (2001) and Milligan (2002) use Canadian micro data to estimate the effect of taxes on household portfolio allocation. Both authors focus on tax-favored retirement savings accounts (Registered Retirement Savings Plans, - RRSPs). Neither exploits the identification strategy that we propose but rather they employ a more traditional approach based on temporal and/or provincial variation in tax rates (following the work on U.S. tax reforms cited above). Veall examines the Canadian tax reform of 1988. He finds a *negative* relationship (though not statistically different from zero in all specifications)

[2] Indeed, one can think of our empirical strategy as a natural experiment or difference - in – difference approach, but where the contrast is across household types (rather than before and after a tax reform), and where the control groups are drawn from a otherwise-similar jurisdiction with a different tax system (rather than from individuals in the same jurisdiction who were less affected by the reform.) The strategy rests on the proposition that the U.S. and Canada are sufficiently similar in preferences, technology and institutions to allow the identification of policy effects. This is an idea with a substantial pedigree (see for example, Card and Freeman, 1993). The contrast between Canada's system of individual taxation and the U.S. system of joint taxation has recently been exploited by Schuetze (2006) to study income splitting among the self-employed.

between RRSP contributions and marginal tax rates – contradicting the prediction that less taxed households face a smaller incentive to invest in tax-favored assets. In contrast, Milligan, who looks at participation (rather than contributions conditional on participation) and uses a combination of temporal and cross-province variation in tax rates, finds that that a 10 percentage-point increase in the marginal tax rate increases the participation probability by eight percent. Milligan notes that a potential explanation for Veall's finding is that tax changes in the period he studied may be overwhelmed by general trends in RRSP behavior. This is one example of the kind of concern with the traditional identification strategy which we outlined above. In contrast, our proposed identification strategy does not employ temporal variation. To the best of our knowledge, our strategy has not been previously employed (neither in Canada nor other jurisdictions).

To preview our main results, we find that Canadian households do shift capital income within the household to take advantage of the system of individual taxation; and that in Canada, holding wealth and total income constant, households with more equal income shares (and hence a higher MTR) tilt their portfolios towards less taxed assets. Moreover we find no evidence of the latter effect in the SCF or PSID, suggesting that the effect we observe in the Canadian data is a true tax response, and not attributable to heterogeneity in preferences or intra-household bargaining. We find that a ten percentage point increase in marginal tax rates increases the mean portfolio share of taxed-favored assets by 2 percentage points, a modest, but statistically significant effect.

The Section 2 presents some evidence that Canadian couples allocate financial assets among partners in order to minimize tax liabilities. We then turn to the effect of

marginal tax rates on portfolio choice. Section 3 elaborates on our data and methods, and results are presented in Section 4. Section 5 concludes.

2. The Allocation of Investment Income Within Households

Variation in the distribution of labor income within households with the same total income generates variation in the effective MTR on capital income if households allocate capital income across household members in order to minimize their tax liability. Typically this would mean having the partner with lower labor income hold financial assets which generate taxable income. Stephens and Ward-Batts (2004) report evidence that U.K. households follow such a strategy. Their study is based on the switch from joint to individual taxation in the U.K. in 1990. Under joint taxation, the MTR on capital income did not depend on the allocation of capital income among partners, so the shift to individual taxation created an opportunity for couples to avoid taxation by attributing capital income to the lower earning partner (usually the wife). Stephens and Ward-Batts report a significant increase in the share and incidence of capital income claimed by wives. They also note an increase in the incidence of couples in which all capital income is attributed to the wife. They therefore conclude that couples responded to the switch from joint to individual taxation by reallocating asset ownership to the partner with lower labor income.

While the Canadian and post-1990 UK tax systems have individual taxation in common, they differ in ways that may be important for our study. In the U.K., couples can choose the division of household assets between spouses. In Canada, this is less

straightforward. In particular, transfers of ownership between couples could attract taxation. Nevertheless, there are likely ways in which Canadian couples can arrange the attribution of capital income so as to minimize tax liability. We now present some evidence that this is case.

The Canadian income tax system is generally based on individual taxation, but particular features of the tax code have at times created incentives similar to those provided by a system of joint taxation. The Canadian tax reform of 1988 is particularly interesting because it reduced the "jointness" of the tax system facing couples in Canada. It eliminated a connection between a secondary earners' effective marginal tax rate (on labor or capital income) and her (or his) spouse's marginal income tax rate. A spousal exemption (tax deduction) was replaced with a non-refundable tax credit. Under both a deduction and a credit, the amount that can be claimed by the primary earner is in each case reduced as the secondary earner's income rises. However, because a deduction reduces taxable income, its value depends on the marginal tax rate of the primary earner who claims the deduction and is therefore much higher for high-income (and hence high marginal tax rate) husbands. In contrast, the value of the tax credit does not depend on the primary earner's marginal tax rate. Thus the effect of the reform was to significantly reduce the "first dollar" marginal tax rate of women married to high-income husbands, while leaving the "first dollar" marginal tax rate of women married to lower income husbands essentially unchanged.

Crossley and Jeon (2007) exploit this change in a difference-in-difference framework to study the effects of taxes on the labor supply of married women. They use

data from the Canadian Survey of Consumer Finances (SCF) for the years from 1986 to 1991[3] and focus on low education women (because these women are most likely to be secondary earners.) Crossley and Jeon report that low education women married to higher income husbands significantly increased their labor force participation (particularly part-time participation) as a result of the Canadian federal tax reform in 1988.[4]

We have used the methods and data of Crossley and Jeon (2007) to study the effect of the 1988 tax reform on the capital income reported by low education married women and their husbands.[5] The idea is that for some of these women (those married to high earning husbands) there was a significant decrease in their 'first dollar' effective marginal tax rate, whether that dollar was labor income or capital income. Crossley and Jeon have shown that some households responded by the wife entering the labor force; here we ask whether some of these households responded by reallocating asset ownership to take advantage of the possibilities for tax saving that the reform introduced.

Table 1 suggests that this is indeed the case. Relative to a control group that was unaffected by the tax reform, wives that experienced a significant decrease in their effective marginal tax rate were 8.5% points more likely to report capital income after the reform. The estimated effect on dollars of capital income reported by these married women is $209 which implies that the treated group essentially doubled their reported capital income. These results echo the findings of Stephens and Ward-Batts (2004). On this basis, we conclude that (like their U.K. counterparts) Canadian couples reallocate

[3] The Canadian SCF is quite similar to the US March supplement to the CPS
[4] The estimated effect on participation rates is sizeable: 9 to 10 percentage points.
[5] An interested reader can find further details on the tax reform, data and methods in Crossley and Jeon (2007).

asset ownership to minimize tax liabilities. This means that the effective marginal tax rate on capital income is often the marginal tax rate of the partner with lower labor income, and so within couples with the *same* household labor income, marginal tax rates on capital income will vary depending how much of that labor income is earned by each partner. The rest of our analysis employs this source of variation to study the effects of taxation on portfolio choice.

3. Portfolio Choice: Data and Methods

3.1 Data

Our main estimates are based on master files from the Canadian Survey of Financial Security, SFS. This survey involved personal interviews in May and June of 1999. The sample includes a supplement of 2,000 households selected from geographical areas with a larger concentration of high income households. Sample weights provided by the survey are used to make the data representative of the Canadian population as a whole.

The SFS (1999) individual files contain information on labor income of all members in a household separately. The portfolio allocation information is recorded at the household level and it is available through the family files. We first merge the SFS individual files with the family files. We then categorize the assets reported in the SFS by their tax characteristics (details below). We define individual income, as the sum of

wages and salaries, self employment income (business and farms), pensions and taxable government transfers.[6]

For the placebo tests we use two major American data sets; the Survey of Consumer Finance, SCF (1998) and the Panel Study of Income Dynamics, PSID (1999). The SCF is a triennial survey that collects information from approximately 4500 respondents concerning household wealth and its allocation. The survey is considered to be the best source of information on household finances in the United States. The main shortcoming of the SCF for our purposes is that detailed income information is not available at the individual level. Thus we must construct within household incomes shares from data on wages and salaries only. (In contrast, in the Canadian SFS, we construct incomes shares that reflect all of the income sources listed in the previous paragraph.) To ensure that this difference is not driving the results, we repeat our placebo tests on the PSID (which has complete income information at the individual level.)

The PSID is a long running panel survey (since 1968) with detailed individual income information. Wealth supplements were added to the main survey every 5 years beginning in 1983. The portfolio allocation information in these supplements is not as detailed as in the SCF. Nevertheless, the aggregate asset categories we need for comparison to the Canada data can be constructed. In the PSID we define income (at both the individual and household level) as the sum of wages and salaries (including overtime payments, bonuses and commissions), child support, and government transfers.

[6] Government transfers include Old Age Security, Canada/Quebec Pension Plan Benefits, disability and death benefits, child benefit and employment insurance benefits. We do not include investment income, because it is clearly endogenous to the shares.

Tax characteristics of different savings instruments depend not only on the type of assets held but also where those assets are held. For example, dividend payments are taxed at the household's marginal tax rate on labor income if stocks are held directly or in mutual funds, but they are not taxed until withdrawn if the stocks are held in a tax-deferred account.

We classify interest bearing assets as heavily taxed assets as the income generated by these assets are taxed at the household's marginal tax rate on labor income. We classify stocks and mutual funds as moderately taxed assets as capital gains are not taxed until realization and generally, households' marginal capital gain tax rates are lower than ordinary income marginal tax rates. One problem we face is that dividend payments are treated similarly to interest earnings and taxed at the household's marginal tax rate on labor income. By classifying stocks as moderately taxed we are assuming that capital gains are the most important part of the returns generated by stocks and dividend payments are relatively less important for tax considerations. Other assets that we classify as moderately taxed include tax-free bonds and tax-free bond funds as capital gains on these assets are taxed even though the interest income they generate is tax-exempt. Our final category includes assets that are tax-favored (deferred) such as RRSPs (locked in retirement accounts) and educational saving accounts in Canada and IRA and Keogh accounts in the US. Interest income, dividend payments, capital gains in these accounts are not taxed until withdrawn. Table 2 summarizes the definitions of asset categories we constructed for the SFS, the SCF and the PSID.

Table 3 presents mean portfolio shares of all three assets categories in total financial wealth and mean income share of minor income earner (in total household income) across the three data sets. While the distribution of income shares of the minor income earner is very similar in both countries (with a mean of about 22% and a standard deviation of 18%), portfolio holdings are quite different in the U.S. and Canada. In particular, tax-favored investment accounts are much more important in Canada than in the United States. The difference is almost entirely offset by the holdings of heavily taxed assets which are much more important in the US. In terms of moderately taxed assets (stocks and mutual funds), the two countries are very similar.[7]

Our working sample includes married (or common-law) couples with or without children[8]. We eliminate households who report negative total income for the survey year. We define the household head as the major income earner of the family (in cases where both spouses earn the same amount, we treat the older one as the head). Households whose heads are older than 65 or younger than 25 are excluded. Finally, households whose heads are full time students during the survey year are eliminated. The final estimation sample size is 6010 households in the SFS. From the sample, for some estimates, we exclude households that are Quebec residents and households with self employment income (because we cannot calculate marginal tax rates for these households). This leaves a working sample of 3710. Identical sample restrictions are applied to the SCF and PSID data, resulting in samples of 1837and 2097 couples.

[7] There are a small number of households that report no holdings of financial assets. We follow Poterba and Samwick (2002) and assume that these households have under-reported holdings of chequing/saving accounts. We therefore assume that these households are censored at 1 for heavily taxed asset group and 0 for other 2 groups. We get very similar results if we simply delete these households.

[8] We exclude the multi-family households from our analysis.

3.2 Methods

We begin our analysis with cross sectional 'reduced form' regressions of portfolio shares $\left(s_h^k\right)$ on household characteristics including household income and wealth (X_h), and the share of household labor income earned by the lower earning partner (ES_h):

$$s_h^k = X_h\beta^k + \alpha^k ES_h + e_h^k \quad \textbf{(1)}$$

Note that k indexes asset classes and h indexes households; e_h^k is a regression disturbance. The variable ES_h is our source of variation in marginal tax rates. We do this on samples that include and exclude the province of Quebec, and include and exclude households with significant self-employment income. Household wealth and income variables are entered as quartile dummies. Additional control variables (also included in X_h) are basic demographic attributes of household head, such as age, gender, education and occupation. We also include a dummy variable for home ownership and a dummy variable indicating that the household has a child. A statistical issue arises from the fact that portfolio shares are bounded between 0 and 1. In particular, a significant number of households have shares of heavily taxed assets equal to one (and hence 0 shares for other asset classes.) To address this, we employ two-limit Tobit estimation (with upper and lower limits at 1 and 0).

This empirical strategy requires that ES_h is correlated with marginal tax rates, and that ES_h can be otherwise excluded from the portfolio share equation. To determine

whether ES_h provides significant variation marginal tax rates, we estimate 'first stage' OLS regressions of the form:

$$mtr_h = X_h\theta + \gamma ES_h + \varepsilon_h^k \quad \textbf{(2)}$$

where mtr_h is the marginal tax rate the household faces on its first dollar of capital income. This variable is calculated using the tax simulation program described in Crossley and Jeon (2007). That program is not capable of generating marginal tax rates for households from the province of Quebec (which has a somewhat different tax system from the rest of Canada) or for households with significant self-employment income. Thus we can only perform this regression on the sample that excludes Quebec residents and self-employed households.

To determine whether ES_h is a plausibly exogenous source of variation in marginal tax rates (that is, whether ES_h can be otherwise excluded from the portfolio share equation), we conduct a 'placebo' test. In particular, we re-estimate the portfolio share equations (1) on similar samples drawn from the 1998 Survey of Consumer of Finances (SCF) and the 1999 wealth module of the Panel Study of Income Dynamics (PSID). Because the U.S. has joint taxation, the first dollar MTR on capital income is unaffected by the distribution of labor income within the household. Thus if the effect ES_h on portfolio allocations operates only through its effect on taxes, it should have no effect in the U.S. Alternatively, a correlation between labor income shares and portfolio allocation in these data would suggest important heterogeneity in preferences or household bargaining.

Finally, reduced form equations like (1) are insufficient for policy analysis that requires magnitudes. Thus we estimate equations that directly capture the relationship between the tax rates on capital income that households face and their portfolio allocations:

$$s_h^k = X_h \pi^k + \phi^k mtr_h + u_h^k \quad (3)$$

The effect of marginal tax rates on portfolio allocations is captured by the ϕ^k parameters, all notation is as above, and u_h^k is regression disturbance. We estimate these Tobit equations with an endogenous regressor by full maximum likelihood Tobit, using ES_h as the instrument for the marginal tax rate. For completeness and comparison, we also report Tobit estimates in which we treat the marginal tax rate as exogenous. Again because of the limitations of the tax calculator at our disposal we can only estimate these equations on the sample that excludes residents of Quebec and households with significant self-employment income.

In a set of (exhaustive) share equations, a theoretical restriction is that the marginal effects of any one explanatory variable should sum to zero across the equations. With exhaustive shares and a common set of explanatory variables, linear regression automatically imposes this restriction. However, this is not the case with Tobit estimation. Imposing the "adding-up" restriction on a system of Tobit equations is very cumbersome (see Poterba and Samwick, 2002, and Rosen and Wu, 2004). We therefore follow Rosen and Wu (2004) and first perform unconstrained estimation and then calculate marginal effects and check that the adding up constraint is satisfied. In our data

the marginal effects come very close to summing to zero (and are not statistically different from zero), so we simply report the unconstrained estimates. We now turn to our results.

4. Portfolio Choice: Results

4.1 Reduced form estimates

The left-hand column of Table 4 presents estimates of equation (1), for all three asset categories in the (Canadian) Survey of Financial Security. Coefficients on the income share (α^k) are reported; full results are presented in Appendix Tables A1-A2. These estimates indicate that households in which the income share of the minor earner is higher (that is, closer to 0.5) hold larger portfolio shares in tax-favored assets and smaller portfolio shares in heavily taxed assets. As such households will face higher effective marginal tax rates, this result accords with the prediction that higher taxed households will face greater incentives to invest in tax-favored assets.

As noted in Section 3, some of our subsequent estimates can only be preformed on a restricted sample that excludes the self-employed and residents of Quebec (this is the case whenever we use the Marginal Tax Rate variable). The right-hand column of Table 4 indicates that our basic results holds (although slightly weaker) for this restricted sample. In particular, in the restricted sample, just as in the full sample, households with more equal income shares hold a larger portfolio share in tax-favored assets.

In addition, the full results reported in Appendix Tables A1-A2 indicate that: 1) higher levels of net wealth and income are associated with higher shares of tax-favored

assets and lower shares of heavily taxed assets; 2) Higher education is associated with lower shares of heavily taxed assets and higher shares of the other asset categories (particularly moderately taxed assets – which includes stocks); and 3) having children and owning a house are positively associated with higher shares of tax-favored assets.

4.2 First stage estimates and placebo test

Our interpretation of the results in Table 4 rests on two assumptions: that income share of the minor earner is a significant determinant of marginal tax rates (instrument relevance) and that (conditional on other controls) the income share of the minor earner does not affect portfolio shares except through the marginal tax rate (instrument validity). We examine the first of these assumptions in Table 5. Table 5 reports the coefficient on a regression of marginal tax rates on the income share of the minor earner and our other control variables (equation (2) in the previous Section). These estimates confirm that, controlling for household income, wealth and demographic characteristics, the income share of the minor earner is a very significant determinant of the effective marginal tax rate on capital income faced by the household.

Table 6 takes up the question of instrument validity, via a placebo test. Here we re-estimate Equation 1, for each asset category, on U.S. data drawn from the PSID and SCF. Again, because the U.S. has joint-taxation, the effective household marginal tax rate on capital income should be independent of the within household distribution of labor income (holding total household income and wealth constant). Table 6 shows that that the coefficient on the income share of the minor earner is never statistically different from zero in these estimates. This result is robust to the exclusion of the self-employed from

these samples (full results available from the authors.) This finding strengthens our confidence in the assumption that, in Canada, the income share of the minor earner affects portfolio choice through the effective household marginal tax rate, and not through some other channel (and thus the income share is a valid instrument for the marginal tax rate).

4.3 IV estimates

Finally, Table 7 presents estimates of equation (3), which relates portfolio shares directly to the effective household marginal tax rate on the first dollar of capital income. The estimates in the right-hand column treat the marginal tax rate as endogenous and use the income share of the minor earner as an instrumental variable. For completeness we report, in the left-hand column, estimates that treat the marginal tax rate as exogenous. Because these estimates obviously require the marginal tax rate variable, they were obtained from the sample which excludes Quebec and the self-employed.

The preferred estimates in the right-hand column indicate that higher marginal tax rates are associated with larger portfolio shares for tax-favored assets. This result is strongly statistically significant and in accord with the prediction that households facing high marginal tax rates have greater incentives to hold wealth in tax-favored assets. Interestingly, we do not find this result when we fail to instrument for the marginal tax rate and instead treat it as exogenous (in the left-hand column).

Because these estimates are from Tobit models the coefficient on the marginal tax rate is not a marginal effect (i.e., the coefficient is not the derivative of the expected portfolio share with respect to the marginal tax rate). However, we have calculated the

marginal effect of the marginal tax rate at the mean of the data and these are also reported at Table 7.We find that a ten percentage point increase in marginal tax rates increases the mean portfolio share of taxed-favored assets by 2 percentage points (or, since the mean share of tax-favored assets is about 0.5, about 4 percent.) This is a modest effect.

5. Conclusion

Identifying the effect of taxation on portfolio choice requires plausibly exogenous variation in marginal tax rates. In progressive tax systems, taxes vary with income levels, but wealth or income almost surely affects portfolio choice directly. In systems of individual taxation – like Canada's – couples with the same level of household income (and wealth) can face different effective tax rates on capital income if labor income is distributed differently within households. In this paper, we employ this source of variation to estimate the effect of taxes on portfolio choices, while controlling for household income and wealth. We find statistically significant but economically modest responses to differential taxation. In a placebo test, using data from the U.S. (which has joint taxation), we find no effect of the intra-household distribution of labor income on portfolio choice. The results of this test support the validity of our empirical strategy.

References

Bertaurt, C. C., (1998), "Stockholding Behavior of U.S. Households: Evidence from the 1983-1989 Survey of Consumer Finances" *Review of Economics and Statistics*, 80:263-275.

Card, D. and R. B. Freeman, Editors, (1993), *Small Differences That Matter: Labor Markets and Income Maintenance in Canada and the United States*, University of Chicago Press for the National Bureau of Economic Research.

Crossley, T. F. and S. Jeon, (2007), "Joint Taxation and the Labor Supply of Married Women: Evidence from the Canadian Tax Reform of 1988," *Fiscal Studies*, 28(3):343-365.

Feldstein, M. S, 1976. "Personal Taxation and Portfolio Composition: An Econometric Analysis", *Econometrica*, vol. 44(4), 631-50

Guiso, L., M. Haliassos and T. Japelli, (2002), *Household Portfolios*, MIT Press.

Hubbard, G. R. (1985), "Personal Taxation, Pension Wealth, and Portfolio Composition" *Review of Economic and Statistics*, 67, 53-60

King, M.A. and J.I.Leape, (1998), "Wealth and Portfolio Composition: Theory and Evidence", *Journal of Public Economics*, 69, 155-193

Milligan, K., (2002), "Tax-Preferred Savings Accounts and Marginal Tax Rates: Evidence on RRSP Participation". *Canadian Journal of Economics*, 35(3):436-456.

Poterba, J., (2001), "Taxation, risk taking and household portfolio behavior", *Handbook of Public Economics*, Vol:3 Elsevier, Amsterdam.

Poterba, J., and S., Andrew, (2003), "Taxation and household portfolio composition: US evidence from the 1980s and 1990s", *Journal of Public Economic*, 87, 5-38.

Rosen, H. S. and S. Wu,, (2004), "Portfolio choice and health status," *Journal of Financial Economics*, 72(3), 457-484.

Samwick , A., (2000) "Portfolio Responses to taxation: Evidence from the end of the rainbow" , J. Slemrod (ed.), Does Atlas Shrug, Harvard University Press, Cambridge, MA.

Scholz, John K.(1994) "Tax Progressivity and Household Portfolios: Descriptive Evidence from the Surveys of Consumer Finances," Tax Progressivity and Income Inequality, J. Slemrod (ed.), Cambridge University Press, 219-267

Schuetze, H. J., (2006), "Income Splitting Among the Self-Employed," *Canadian Journal of Economics*, 39(4):1195–1220.

Slemrod, J., (2002), "Tax avoidance, evasion and administration", *Handbook of Public Economics*, Vol:3 Elsevier, Amsterdam

Stephens, M. Jr. and J.Ward-Batts, (2004), "The impact of separate taxation on the intra-household allocation of assets: evidence from the UK," *Journal of Public Economics*, 88:1989– 2007

Veall, M. R. (2001) "Did tax flattening affect RRSP contributions?" *Canadian Journal of Economics* 34, 120-31

Tables

Table 1: Effect of the 1998 Canadian Tax Reform on Capital Income Reported by Low Education Women Married to Higher income men (Simple Diff-in-Diff estimates)

a) Incidence of Capital Income (%)

Group	Pre tax reform	Post tax reform	Difference	Difference in Difference
Control (married to low-income husband)	15.1	18.5	3.4	
Treatment (married to high-income husband)	19.8	31.7	11.9	8.5** (2.9)

b) Dollars of Capital Income

Group	Pre tax reform	Post tax reform	Difference	Difference in Difference
Control (married to low-income husband)	119	227	108	
Treatment (married to high-income husband)	202	519	317	209** (84.4)

Notes to Table 1:

1. Standard errors in parentheses.
2. ** statistically significant at the 5% level.

Table 2: Asset Classification

	Heavily Taxed Assets	Moderately Taxed Assets	Tax-Favored
SFS (1999) Canadian	a) Bonds (Saving + Other) b) Term Deposits c) Guaranteed Income Certificates d) Mortgage Backed Security Funds e) Cheq. & Saving Accounts f)T-bills	a) Non-RRSP Stocks b) Mutual funds and other investment funds exclusive of RRSP c)Trust funds	a) Registered educational savings b) RRSPs c) Home ownership savings plan funds
PSID (1999)	a)Checking & Savings accounts b) Money market funds c)Certificates of deposit d)government savings bonds e)T-bills f)Bond funds, cash value in a life insurance policy, trusts	Directly held publicly and privately issued stocks and mutual funds	IRA and Keogh accounts
SCF (1998)	a) Cheq. & Saving Accounts b) Money market funds (excluding tax-free ones) c) CDs d) Savings bonds e) Mortgage-backed bonds f) Corporate Bonds g) Foreign Bonds h) Cash or call money accounts i)T-bills j)government bond funds and other bond funds	a) Stock mutual funds b)Stocks c)trusts d)Tax-free Money market funds e) Tax-free bonds f) Tax-Free Bond Funds	IRA and Keogh accounts

Table 3: Descriptive Statistics

AVERAGE PORTFOLIO & INCOME SHARES

	SFS(1999)		SCF(1998)	PSID(1999)
	Full	**Restricted***		
Income Share of Lower Earning Partner	.236 (.178)	.237 (.179)	.210 (.188)	.225 (.180)
	Portfolio Shares			
Heavily Taxed	.389 [.412]	.388 [.411]	.632 [.632]	.671 [.683]
Moderately	.081 [.304]	.080 [.301]	.208 [.528]	.139 [.459]
Taxed-Favored	.530 [.697]	.531 [.705]	.161 [.429]	.190 [.515]

Notes to Table 3:

1. For Income Shares standard deviations are reported in round parentheses ()
2. For Portfolio Shares conditional Averages are reported in square parentheses [].Conditional average refers to only households who have positive amount of the asset group
3. For SFS and SCF survey weights are used in all calculations.
4. *Restricted Sample excludes the self-employees and Quebec residence from the full sample.

Table 4: Reduced Form Tobit Estimates (Summary)

$$s_h^k = X_h\beta^k + \alpha^k ES_h + e_h^k$$

Coefficients on the Income Share(α^k)

	SFS (1999)	
	Full Sample	Restricted Sample
Heavily Taxed	-0.053*	-0.048
	(.03)	(.039)
Moderately Taxed	-0.062	-0.088
	(.048)	(.067)
Tax-Favored	.092***	0.087**
	(.034)	(.044)

Notes to Table 4:

1. Standard Errors are in parentheses.
2. For complete results see Appendix Tables A1-A2.
3. Each tobit model allows for censoring both below and above (at portfolio shares of zero and one.)
4. *** significant at 1%, ** significant at 5%; ** significant at 10%

Table 5: First Stage Estimates (SFS)

$$mtr_h = X_h\theta + \gamma ES_h + \varepsilon_h^k$$

Coefficients on the Income Share(γ)	
Income Share	0.346***
	(0.01)
R-squared	*0.43*

Notes to Table 5:

1. Standard Errors are in parentheses.
2. For complete results see Appendix Table A3
3. *** significant at 1%

Table 6: Placebo Tests (Reduced Forms on SCF and PSID)

$$s_h^k = X_h\beta^k + \alpha^k ES_h + e_h^k$$

Coefficients on the Income Share(α^k)		
	SCF(1998)	**PSID (1999)**
Heavily Taxed	.039	-.034
	(.024)	(.085)
Moderately Taxed	-.046	*
	(.030)	
Tax-Favored	.014	.035
	(.029)	(.10)

Notes to Table 6:

1. Standard Errors are in parentheses.
2. For complete results see Appendix Tables A4-A5.
3. Each Tobit model allows for censoring both below and above (at portfolio shares of zero and one.)
4. Results are qualitatively similar in the PSID and SCF when the self-employed are deleted from the sample (importantly, the income share remains insignificant). Full results are available from authors.
5. * Did not converge.

Table7: IV Estimates (Marginal Tax Rates instrumented by Income Share)

$$s_h^k = X_h \pi^k + \phi^k mtr_h + u_h^k$$

Coefficients on the Marginal Tax Rate (ϕ^k)

SFS-Restricted Sample

	TOBIT		***IVTOBIT***	
	Coef.	Marginal	Coef.	Marginal
Heavily Taxed	-0.029	-.020	0.027	.023
	(.057)		(.074)	
Moderately Taxed	-0.003	-.0007	-0.222	-.046
	(.09)		(.178)	
Tax-Favored	0.058	.040	0.255**	.212
	(.064)		(.114)	

Notes to Table 7:

1. Standard Errors are in parentheses.
2. For complete results see Appendix Tables A6-A7.
3. Each Tobit model allows for censoring below and above (at portfolio shares of zero and one.)
4. *** significant at 1%, ** significant at 5%; ** significant at 10%

Appendix

Table A1: Tobit Estimates of Reduced Forms (Full Sample)

	Share of		
	Heavily Taxed	*Moderately Taxed*	*Tax Favored*
Income Share of Lower Earning Partner	**-0.053**	**-0.062**	**0.092**
	(.030)	**(.048)**	**(.034)**
Income			
2. Quartile	-0.137	-0.002	0.15
	(.014)	(.026)	(.017)
3. Quartile	-0.163	-0.018	0.196
	(.015)	(.026)	(.018)
4. Quartile	-0.161	0.025	0.175
	(.018)	(.028)	(.020)
Net Worth			
2. Quartile	-0.635	0.402	0.622
	(.015)	(.033)	(.017)
3. Quartile	-0.771	0.56	0.737
	(.016)	(.034)	(.018)
4. Quartile	-0.821	0.851	0.676
	(.018)	(.037)	(.020)
Age			
35-44	0.03	-0.08	-0.018
	(.014)	(.022)	(.015)
45-54	0.04	-0.115	-0.026
	(.016)	(.025)	(.017)
55-65	0.039	-0.159	-0.002
	(.018)	(.030)	(.021)
Education			
High School	-0.057	0.073	0.049
	(.017)	(.030)	(.019)
Some College	-0.106	0.131	0.076
	(.022)	(.036)	(.025)

Continues on next page

College Degree	-0.062	0.101	0.043
	(.015)	(.026)	(.017)
Post College	-0.065	0.168	0.003
	(.021)	(.033)	(.024)
Male	-0.018	-0.022	0.03
	(.013)	(.021)	(.015)
Have Children	-0.024	-0.004	0.023
	(.012)	(.019)	(.013)
Owning House	-0.064	0.052	0.057
	(.014)	(.024)	(.016)
Constant	1.238	-0.803	-0.328
	(.026)	(.05)	(.03)

Notes to Table A1:

1. Standard Errors are in parentheses.
2. Each Tobit model allows for censoring below and above (at portfolio shares of zero and one.)
3. The regressions also include indicator variables for the occupation and industry of the head of households. For the brevity of the tables we didn't report them.

Table A2: Tobit Estimates of Reduced Forms (Restricted Sample)

	Share of		
	Heavily Taxed	*Moderately Taxed*	*Tax Favored*
Income Share of Lower Earning Partner	-0.048	-0.088	0.087
	(.039)	(.063)	(.044)
Income			
2. Quartile	-0.136	0.019	0.146
	(.02)	(.036)	(.022)
3. Quartile	-0.173	0.036	0.189
	(.021)	(.037)	(.024)
4. Quartile	-0.148	0.069	0.148
	(.024)	(.039)	(.026)
Net Worth			
2. Quartile	-0.703	0.415	0.706
	(.02)	(.046)	(.023)
3. Quartile	-0.855	0.557	0.849
	(.022)	(.048)	(.025)
4. Quartile	-0.918	0.856	0.803
	(.024)	(.051)	(.027)
Age			
35-44	0.03	-0.091	-0.011
	(.017)	(.028)	(.019)
45-54	0.033	-0.15	-0.012
	(.020)	(.033)	(.022)
55-65	0.035	-0.167	0.001
	(.023)	(.039)	(.026)
Education			
High School	-0.026	0.071	0.014
	(.022)	(.039)	(.025)
Some College	-0.08	0.106	0.061
	(.028)	(.047)	(.032)
College Degree	-0.053	0.068	0.045
	(.019)	(.034)	(.021)
Post College	-0.059	0.147	0.006
	(.028)	(.045)	(.032)

Continues on next page

Male	-0.015	-0.007	0.019
	(.018)	(.028)	(.020)
Have Children	-0.03	-0.009	0.027
	(.015)	(.025)	(.017)
Owning House	-0.032	0.06	0.014
	(.017)	(.031)	(.02)
Constant	1.287	-0.846	-0.377
	(.033)	(.067)	(.038)

Notes to Table A2:

1. Standard Errors are in parentheses.
2. Each Tobit model allows for censoring below and above (at portfolio shares of zero and one.)
3. The regressions also include indicator variables for the occupation and industry of the head of households. For the brevity of the tables we didn't report them.

Table A3 - First Stage Estimates

	Marginal Tax Rate
Income Share of Lower Earning Partner	**0.346**
	(.010)
Income	
2. Quartile	0.020
	(.005)
3. Quartile	0.027
	(.005)
4. Quartile	0.086
	(.006)
Net Worth	
2. Quartile	-0.003
	(.005)
3. Quartile	-0.005
	(.005)
4. Quartile	.021
	(.006)
Age	
35-44	-0.011
	(.004)
45-54	-0.012
	(.005)
55-65	-0.016
	(.006)
Education	
High School	0.002
	(.005)
Some College	0.005
	(.007)
College Degree	0.004
	(.005)
Post College	-.0003
	(.007)

Continues on next page

Male	0.009
	(.004)
Have Children	0.026
	(.004)
Owning House	.007
	(.004)
Constant	.103
	(.008)
R-squared	*0.43*

Notes to Table A3:

1. Standard Errors are in parentheses
2. The regressions also include indicator variables for the occupation and industry of the head of households. For the brevity of the tables we didn't report them.

Table A4: Placebo Test: Tobit Estimates of Reduced Forms

Survey of Consumer Finances (1998)

	Share of		
	Heavily Taxed	*Moderately Taxed*	*Tax Favored*
Income Share of Lower Earning Partner	0.039	-0.046	0.014
	(.024)	(.030)	(.029)
Income			
2. Quartile	-0.009	0.001	0.026
	(.013)	(.018)	(.017)
3. Quartile	0.042	-0.001	-0.042
	(.015)	(.019)	(.018)
4. Quartile	0.049	0.011	-0.054
	(.016)	(.02)	(.019)
Net Worth			
2. Quartile	-0.684	0.629	0.553
	(.015)	(.021)	(.019)
3. Quartile	-1	0.934	0.74
	(.016)	(.023)	(.021)
4. Quartile	-1.122	1.254	0.599
	(.018)	(.025)	(.023)
Age			
35-44	0.01	-0.088	0.085
	(.015)	(.019)	(.019)
45-54	0.003	-0.13	0.126
	(.015)	(.019)	(.019)
55-65	0.021	-0.226	0.197
	(.017)	(.021)	(.021)
Education			
High School	-0.019	-0.025	0.04
	(.014)	(.018)	(.017)
Some College	0.019	-0.054	0.041
	(.017)	(.022)	(.020)
College Degree	-0.059	0.022	0.07
	(.014)	(.018)	(.017)

Continues on next page

Post College	-0.068	0.032	0.084
	(.015)	(.018)	(.018)
Male	0.014	-0.194	0.3
	(.065)	(.076)	(.094)
Have Children	0.011	-0.003	-0.014
	(.01)	(.012)	(.012)
Owning House	-0.103	0.113	0.06
	(.013)	(.017)	(.016)
Constant	1.394	-0.332	-1.05
	(.068)	(.081)	(.098)

Notes to Table A4:

1. Standard Errors are in parentheses.
2. Each Tobit model allows for censoring below and above (at portfolio shares of zero and one.)
3. The regressions also include indicator variables for the occupation and industry of the head of households. For the brevity of the tables we didn't report them.

Table A5: Placebo Test: Tobit Estimates of Reduced Forms, PSID (1999)

	Share of	
	Heavily Taxed	*Tax Favored*
Income Share of Lower Earning Partner	-0.034	0.035
	(.085)	(.1)
Income		
2. Quartile	-0.009	0.054
	(.048)	(.057)
3. Quartile	-0.019	0.048
	(.048)	(.057)
4. Quartile	-0.086	0.069
	(.049)	(.058)
Net Worth		
2. Quartile	-0.372	0.403
	(.077)	(.105)
3. Quartile	-1.025	1.026
	(.077)	(.105)
4. Quartile	-1.376	1.302
	(.08)	(.108)
Age		
35-44	-0.027	0.089
	(.04)	(.049)
45-54	0.018	0.061
	(.042)	(.051)
55-65	-0.056	0.166
	(.054)	(.064)
Male	-0.583	0.365
	(.273)	(.298)
Have Children	-0.03	0.056
	(.033)	(.039)
Owning House	-0.107	0.081
	(.043)	(.052)
Constant	2.575	-1.643
	(.291)	(.325)

Continues on next page

Notes to Table A5:

1. Standard Errors are in parentheses.
2. Each Tobit model allows for censoring below and above (at portfolio shares of zero and one.)
3. In Moderately taxed asset group estimation did not converge

Table A6: Portfolio Shares on Marginal Tax Rates

	Share of		
	Heavily Taxed	*Moderately Taxed*	*Tax Favored*
Marginal Tax Rate	**-0.029**	**-0.003**	**0.058**
	(.057)	**(.09)**	**(.064)**
Income			
2. Quartile	-0.138	0.013	0.149
	(.02)	(.036)	(.022)
3. Quartile	-0.178	0.024	0.196
	(.021)	(.036)	(.023)
4. Quartile	-0.152	0.054	0.155
	(.024)	(.04)	(.027)
Net Worth			
2. Quartile	-0.704	0.415	0.707
	(.02)	(.046)	(.023)
3. Quartile	-0.856	0.556	0.85
	(.022)	(.048)	(.025)
4. Quartile	-0.916	0.86	0.8
	(.024)	(.051)	(.027)
Age			
35-44	0.03	-0.089	-0.012
	(.017)	(.029)	(.019)
45-54	0.034	-0.145	-0.015
	(.02)	(.033)	(.022)
55-65	0.038	-0.16	-0.004
	(.023)	(.039)	(.026)
Education			
High School	-0.026	0.072	0.013
	(.022)	(.039)	(.025)
Some College	-0.078	0.109	0.059
	(.028)	(.047)	(.031)
College Degree	-0.052	0.071	0.043
	(.019)	(.034)	(.021)
Post College	-0.056	0.154	0.001
	(.028)	(.045)	(.032)

Continues on next page

Male	-0.011	0.003	0.011
	(.017)	(.027)	(.019)
Have Children	-0.028	-0.005	0.022
	(.015)	(.025)	(.017)
Owning House	-0.032	0.059	0.014
	(.017)	(.032)	(.02)
Constant	1.279	-0.871	-0.364
	(.034)	(.067)	(.038)

Notes to Table A6:

1. Standard Errors are in parentheses.
2. Each Tobit model allows for censoring below and above (at portfolio shares of zero and one.)
3. The regressions also include indicator variables for the occupation and industry of the head of households. For the brevity of the tables we didn't report them.

Table A7: IVTobit Estimates - Marginal Tax Rates instrumented by Income Share

	Share of		
	Heavily Taxed	*Moderately Taxed*	*Tax Favored*
Marginal Tax Rate	**0.027**	**-0.222**	**0.255**
	(.074)	**(.178)**	**(.114)**
Income			
2. Quartile	-0.062	0.021	0.125
	(.014)	(.035)	(.021)
3. Quartile	-0.071	0.041	0.165
	(.015)	(.037)	(.022)
4. Quartile	-0.078	0.086	0.124
	(.019)	(.045)	(.028)
Net Worth			
2. Quartile	-0.104	0.422	0.696
	(.016)	(.046)	(.021)
3. Quartile	-0.171	0.563	0.853
	(.016)	(.047)	(.022)
4. Quartile	-0.206	0.86	0.81
	(.017)	(.051)	(.024)
Age			
35-44	-0.01	-0.093	-0.012
	(.011)	0.028	0.018
45-54	-0.018	-0.152	-0.013
	(.013)	0.032	0.02
55-65	-0.019	-0.159	0.01
	(.016)	0.038	0.024
Education			
High School	0.004	0.07	0.006
	(.015)	(.037)	(.022)
Some College	-0.014	0.113	0.055
	(.019)	(.045)	(.028)
College Degree	-0.016	0.074	0.051
	(.013)	(.033)	(.019)

Continues on next page

Post College	-0.023	0.145	-0.003
	(.018)	(.043)	(.029)
Male	-0.005	0	0.02
	(.011)	(.027)	(.018)
Have Children	0	0.055	0.009
	(.012)	(.03)	(.018)
Owning House	-0.034	0.006	0.012
	(.01)	(.024)	(.015)
Constant	0.445	-0.822	-0.412
	(.027)	(.072)	(.04)

Notes to Table A7:

1. We estimate Tobit Equations with an endogenous regressor by full maximum likelihood Tobit.
2. Income share of minor income earner is used as an instrumental variable.
3. Standard Errors are in parentheses.
4. Each Tobit model allows for censoring below and above (at portfolio shares of zero and one.)
5. The regressions also include indicator variables for the occupation and industry of the head of households. For the brevity of the tables we didn't report them.

V

On the Identification of Intertemporal Allocation Parameters: Instrument and Sampling Issues

1. Introduction

Models based on dynamic optimization problems faced by economic agents have long been the work horses of macroeconomics. Since Hall (1978), first order conditions (known as Euler Equations) from these optimization problems have been used extensively to test these models against economic data and as a basis for estimating preference parameters. This tradition includes: test of the life-cycle model of consumption (see for example, Attanasio et. al. 1999); studies of the investment behaviour of firms (for example, Bond and Meghir 1994, or Mulligan, 2004); and tests of asset pricing theories (for example, Mehra and Prescott, 1985). The main attraction of theEuler equation approach is that it allows researchers to estimate preference parameters with limited data and without fully specifying the stochastic processes that agents face. In principle, researchers do not need to model agents' expectations (retirement expectations, fertility plans etc.) or observe their wealth when estimating preference parameters. Unfortunately these advantages of the Euler Equation approach are significantly diminished by practical problems that arise from the nature of the available data.

Euler equations represent the behavior of an individual agent (households and firms). There is substantial evidence that estimating these models with aggregate data can lead not only to biased parameter estimates but also to false rejections of the underlying models (see for example, Attanasio and Weber, 1993). As a result, researchers have

moved towards using household/firm level panel survey data instead of aggregate time series data. However, this solution generates a further problem because measurement error is endemic to such survey data. This measurement error interacts with the nonlinear structure of the Euler equation: in the presence of measurement error, standard non-linear GMM methods yield inconsistent estimates (Amemiya 1986). This problem is in turn normally addressed by taking a first-, or possibly second-order approximation to the Euler equation (to give a linearized or approximate Euler equation) and then using standard linear IV and GMM techniques to deal with the measurement error.[1]

However, this particular cure may be worse than the disease. Higher order terms that are ignored in the approximation are potentially correlated with the typical instruments (lagged variables) used in these estimation procedures, and this can lead to substantial bias. Consequently, the usefulness of approximate Euler Equations is now very much in question.

Several papers have explored this issue in the context of the consumption Euler equation. The common approach of these studies is that they solve and simulate a life-cycle consumption model and then perform Monte-Carlo experiments with the simulated data to investigate whether a linearized Euler equation yields good estimates of preference parameters (that is, they investigate whether it is possible to recover the preference parameters values that were assumed in solving the model and generating the simulated data). Ludvigson and Paxson (2001) investigate the estimation of the relative

[1] There are alternatives. Alan, Attanasio and Browning (2007) introduce two new nonlinear GMM estimators that deal with measurement error in particular circumstances. Alan and Browning (2006) propose an approach that is not based on Euler equations but rather on modeling expectation errors directly.

prudence parameter in an environment with a fixed interest rate and impatient agents. Following the empirical strategy of Dynan (1993), they employ a second-order approximation to the Euler equation as a basis for estimation. They conclude that this strategy for estimating the prudence parameter is not useful because conventional instruments are correlated with the approximation error. Carroll (2001) reaches the same conclusion for the elasticity of intertemporal substitution (EIS) in an environment with cross-sectional variation in interest rates. A common feature of these studies is the lack of time series variation in interest rates. Attanasio and Low (2004) argue that with sufficiently long sample periods and enough time-series variation in the intertemporal price (interest rate), consistent estimates of EIS can be obtained with linearized Euler equations.

In this paper, we revisit this debate and offer a reconciliation of the differing conclusions reached by different researchers. First, we argue that the studies cited above reached different conclusions in large part because their assumed economic environments differed in ways that were critical to the question at hand. We show that there are environments (parameter values) where linearized Euler equations work well and environments where the instrument problems described above plague estimation. In particular, we show that problems with instrument validity are related to the severity of the approximation bias; that the severity of approximation bias is determined by the curvature of the underlying policy functions; and that the curvature of the policy functions follows from assumptions made in specifying the economic environment. Second, we argue that in choosing between linear - and higher-order approximations,

researchers face a difficult trade-off between instrument validity and instrument relevance.

We proceed as follows. We consider a set up with standard (CRRA) preferences, (aggregate) interest rate uncertainty, and an uninsurable idiosyncratic income risk. We first solve numerically for policy functions in 6 different economic environments (that is, 6 different choices of parameter values), which correspond to the different environments assumed in the literature. Next we develop a simple and intuitive non-parametric statistic to summarize the important characteristic of these different environments. This non-parametric statistic, which we call the degree of "effective curvature", weights the curvature of the policy function (consumption function) by the agents' ex-post cash-on-hand (asset) distribution. While the consumption function varies with different values of preference and income process parameters, it will typically exhibit the same general shape: highly curved at the low wealth levels and near linear at the high wealth levels. Thus the key consequence of different parameter values (especially discount rates) is not so much the shape of the resulting policy function but rather which part of the policy function is ex-post relevant for agents. This is why we develop our measure of effective curvature. Our proposed measure integrates over the ex-post distribution of the key state variable (cash-on-hand) and so captures both the curvature of the policy function per se, and the ex-post relevance of different parts of the state space (and hence different parts of the policy function). We then show that the validity and relevance of instruments typically used in the estimation of approximate Euler equations depend on the degree of the effective curvature, and the differing results of, for example Carroll (2001) and

Attanasio and Low (2004) can be explained by differences in the effective curvature in their environments.

The second contribution of our analysis is to assess not only validity but also instrument relevance. The issue of instrument relevance (or weak instrument problems) has been discussed for consumption Euler equations estimated on aggregate data (Yogo, 2004). However, we are not aware of any analysis of instrument relevance in the case of approximate Euler equations estimated on micro data. In this regard, our simulations demonstrate two things. First, instrument relevance is also linked to the effective curvature of the consumption function, and in particular there is a sharp trade-off between instrument validity and instrument relevance. The instruments typically used to estimate consumption Euler equations tend to be strongly relevant in environments with high effective curvature. These are, of course, exactly the environments where the instruments are not valid. Conversely, in environments with low effective curvature, the typical instruments tend to be less relevant, leading to the possibility of weak instrument problems. Second, we show that the choice between (log-)linear and second-order approximations faces a similar trade-off. Using a second-order approximation reduces problems with instrument validity, but exacerbates weak instrument problems.

A third result that arises from our simulations is that superior estimates of the EIS might be obtained by careful sampling or conditioning on cash-on-hand (wealth). In the final section of the paper we present an empirical illustration using the American Panel Study of Income Dynamics (PSID). The objective of this exercise is to implement the conditioning strategies suggested by our simulations. Our empirical results accord with

the simulation results. By sampling wealthy households (agents that carry forward positive assets for several (lagged) periods) we get sensible and statistically significant estimates of the EIS.

2. The Econometrics of Euler Equation Approximation

The first order condition from a standard intertemporal consumption optimization problem without liquidity constraints is:

$$U'(C_{t-1}) = \beta E_{t-1}[(1+R_t)U'(C_t)] \tag{4}$$

where U' is the marginal utility of consumption, β is the discount factor and R_t is the real rate between periods t-1 and t. A widely used functional form for the sub-utility function is the iso-elastic form:

$$U'(C_t) = (\frac{C_t}{1-\gamma})^{1-\gamma} \tag{5}$$

where the parameter γ is the coefficient of relative risk aversion. Usually interest centers on the reciprocal of this parameter, $(\frac{1}{\gamma})$, the Elasticity of Intertemporal Substitution (EIS). Substituting this utility function into equation (1) yields an exact Euler equation:

$$(\frac{C_t}{C_{t-1}})^{-\gamma}(1+R_t)\beta = \varepsilon_t \text{ with } E_{t-1}(\varepsilon_t) = 1 \tag{6}$$

This relationship has been the basis of very many estimates of the preference parameters (β, γ) and tests for the validity of the standard orthogonality assumptions in general and for the "excess sensitivity" of consumption to predictable income growth in particular. GMM estimation is based on the assumed orthogonality of the error term ε_t to all

variables dated t-1 or before, such as lagged consumption, interest rate and income variables. As originally emphasized by Hall (1978), this is a very attractive procedure since one can estimate the preference parameters without explicitly parameterizing the stochastic environment that agents face.

Problems for GMM estimation on micro data arise if the consumption data are measured with error. For example, if we allow for a multiplicative measurement error so that observed consumption C_t^0 is given by:

$$C_t^0 = C_t \eta_t \tag{7}$$

then the exact Euler equation for observable consumption becomes

$$(\frac{C_t^0}{C_{t-1}^0})^{-\gamma}(1+R_t)\beta = (\frac{\eta_t}{\eta_{t-1}})^{-\gamma}\varepsilon_t \tag{8}$$

The problem is that the composite error term does not have a conditional expectation of unity, even if we assume that η_t and ε_t are independent[2]:

$$E_{t-1}[(\frac{\eta_t}{\eta_{t-1}})^{-\gamma}\varepsilon_t] = E_{t-1}(\frac{\eta_t}{\eta_{t-1}})^{-\gamma}\ E_{t-1}(\varepsilon_t) = E_{t-1}(\frac{\eta_t}{\eta_{t-1}})^{-\gamma} \neq 1$$

It is now widely accepted that household level consumption data is likely to be very noisy. For example, Runkle (1991) estimates 76% of the variation in the growth rate of food consumption in the PSID is noise. Dynan (1993) reports the standard deviation of changes in log consumption in the CEX (American Consumer Expenditure Survey) is 0.2, which seems too large for `true' variations. The other widely used data sources are quasi-panels, constructed from cross-section expenditure survey information by taking

[2] Note that η_{t-1} is not in the agent's information set at time t-1 and cannot be taken outside the conditional expectation.

within-period means following the same population (e.g. means over all the 25 year olds in one year and all the 26 year olds in the next year). Although this averaging reduces the effect of measurement error, the construction of quasi-panels from samples which change over time induces sampling error that acts very much like measurement error.

One way to deal with measurement error problem is to linearize the equation (3) and use standard IV and GMM techniques. In particular, the convention is to assume measurement error in consumption is multiplicative. Naturally, log-linearization will move such measurement error into the expectation errors. By using twice lagged instruments (due to the MA(1) structure of the errors created by the measurement error), identification can be achieved Following the steps in Carroll (2001)[3] equation (6) can be derived[4]:

$$\Delta \log C_t = \alpha - \frac{1}{\gamma}\log(1+R_t) + \frac{\gamma+1}{2}(\Delta \log C_t)^2 + e_t \qquad (9)$$

where the constant term α contains the discount rate and the unconditional means of the third (and higher order) moments of the error term ε_t. The error term e_t contains the expectation error and also time varying components of the higher conditional moments (conditional on past information). The first order log-linear approximation (equation (6) without the squared term) has been used extensively in the applied micro literature.

[3] See also Hansen and Sigleton (1983) , Attanasio and Low (2004).
[4] Note that because of using iso-elastic form, one parameter γ controls both EIS (through the coefficient on the real rate) and prudence terms (through the coefficient on the square term). This is a special case for iso-elastic sub-utility form.

2.1 Concavity of Consumption Function and Approximation Bias

The problem with the approximate Euler equation is that the time varying components of higher order moments that are subsumed into the error term will cause it to be correlated with lagged variables, which leaves the researchers without any valid instruments. We refer to the resulting estimation bias as approximation bias. Naturally, this instrument invalidity problem becomes more serious as higher order moments become more prominent in the error term. The significance of these moments is strongly related to the concavity of the underlying policy rules. Consequently, the concavity of the region of the consumption function that is ex-post relevant for the agents (the effective curvature) determines the severity of approximation bias.

Note that our notion of "effective curvature" is not the same as the overall concavity of the consumption function. We know that in all environments consumption functions are highly curved at low wealth levels but near-linear at high wealth levels. What affects the severity of approximation bias is not this curvature but where actually agents operate under these curved consumption functions. For example, since patient agents are fast wealth accumulators, they operate around the near linear part of consumption functions. As a result their effective curvature is low even though their consumption functions' overall curvature is not different from that of patient agents. This is extremely important because as we will demonstrate, instrument validity and relevance depend strongly on this.

We measure the effective curvature by weighting the curvature of underlying consumption function (at a given age) at every point by the ex-post density of normalized

cash-on-hand[5]. In particular to do this, we allow consumption values from the policy function to be a flexible function of the normalized cash-on-hand (labeled as x). That is, first we estimate:

$$c = g(x) + u$$

and then using the estimates we calculate the standard curvature measure of the consumption function $\left(-\frac{c''(x)}{c'(x)}x\right)$ at every x point. At the last step we take the population weighted average of this measure. Thus our nonparametric method calculates:

$$\textit{Effective Average Curvature} = -\int f(x)\frac{c''(x)}{c'(x)}xdx \qquad (10)$$

The effective curvature of the consumption function is connected to instrument validity through the distributional behavior of the expectational errors. As effective curvature increases, higher order moments become more prominent and cause instrument problems.

In order to illustrate our point we examine six different life cycle models of consumption. In all models, within period utility is iso-elastic with the coefficient of relative risk aversion set to 4. We generate 80 periods of consumption paths for ex-ante identical consumers. We then remove the first 20 periods and the last 20 periods. Further details of simulation methods are given in the Appendix. Agents face two types of income shocks, permanent and transitory. The income process of agent h is:

$$Y_{h,t} = P_{h,t}U_{h,t} \qquad (11)$$

[5] This is the ratio of cash on hand to permanent income.

where $U_{h,t}$ is an iid lognormal transitory shock with unit mean and a constant variance $\left(e^{\sigma_u^2}-1\right)$ and $P_{h,t}$ is permanent income which follows a log random walk process:

$$P_{h,t} = GP_{h,t-1}Z_{h,t} \tag{12}$$

where $Z_{h,t}$ is an iid lognormal permanent shock with unit mean and a constant variance $\left(e^{\sigma_z^2}-1\right)$. In our simulations we set σ_u^2 to 0.1 and σ_z^2 to 0.05; these values are in line with those used in the literature and experiments with other values give qualitatively similar results. [6] We assume that the innovations to income are independent over time and across individuals so that we abstract aggregate shocks to income. However, aggregate shocks are not completely absent since the realizations of the real rate are assumed to be the same across agents. The real rate follows an AR(1) process with a mean of 0.03, an AR parameter of 0.6 and 0.025 for the standard deviation of error.

Table 1 presents all the parameter values assumed for the 6 models. The models differ by degree of impatience and by the presence or absence of a borrowing constraint. Table 2 summarizes all six models' distinguishing features. Model 1 is similar to the environment studied by Attanasio and Low (2004). Agents' discount rates are equal to the mean real interest rate (0.03) and there is no borrowing constraint (except for the life-time resource constraint). The important feature of this model is that even though borrowing is allowed up to the natural limit, individuals do not borrow because they are quite patient and have a strong taste for accumulation. The second model is the impatient

[6] See Alan and Browning (2006)

version[7] of the first model. In this model the discount rate of agents set to 0.07. As a result, agents borrow especially early in their life cycle.

Model 3 and Model 4 are used by Carroll (2001). In these models transitory income shocks, with some small probability (0.01), can take a '0' value in any period. With this addition to the model, the precautionary motive of agents is heightened. This assumption, along with concavity of the consumption function and backward induction imposes an implicit borrowing constraint. The resulting consumption functions are highly curved at low wealth levels. Model 4 is the impatient version of model 3.

Finally, we examine two models that were first proposed by Deaton (1991) but not previously considered in the literature on Euler equation estimation. In these models, individuals are explicitly prevented from borrowing. This assumption (with a lower bound for labor income) leads to a kink in the consumption function. We have two motivations for including these models. First, the patient version (Model 5) of the Deaton environment is very similar to Model 1's environment (where borrowing is allowed). This helps to clarify our argument about effective (not overall) curvature, and instrument problems. In addition, we believe that the Model 6 (Deaton impatient) resembles the actual buffer stock behavior much more accurately than any of our models[8].

Figure 1 summarizes the general characteristics of our models. For each model, the consumption function and distribution of normalized cash-on hand at age 40 are

[7] Since, without loss of generality, we assume away predictable income growth, degree of impatience is measured by the difference between real interest rate and individual's discount rate.

[8] Note that in this model a liquidity constraint often binds and when this occurs Euler equation does not hold. Therefore if the agent does not carry forward assets between period t-1 to t, we remove those two periods.

plotted; also, effective curvatures are reported in the top right corner of graphs. It is clear from the figures that cash-on-hand distributions of patient models (models 1, 3 and 5) are located at higher wealth levels. As a result, for individuals in these environments curved parts of the consumption function are irrelevant and their effective curvature is low. On the contrary, in impatient models (models 2, 4 and 6) agents accumulate very little wealth. They act like buffer stock savers, and mostly operate on the curved part of their consumption function. As a result the effective curvatures for these models are much higher than their patient counterparts.

Table 3 shows the second, third and fourth moments of expectation errors in these different environments. This table exhibits the link between effective curvature and approximation bias through the distributional behavior of expectation errors. Clearly, as the effective curvature increases higher order moments become more prominent. Model 6 is the only exception to this; in the following section we explain this departure from the pattern.

2.2 Effective Curvature, Validity and Relevance of Instruments

In this subsection we investigate the effect of effective curvature on instrument validity and relevance in the estimation of approximate Euler equations. To examine instrument validity, we first construct the true composite errors of the first and second order log-linearized models using the true parameter values. That is

$$\Delta \log C_t + \frac{1}{4}\log(1+R_t) = e_t \tag{13}$$

for the first order approximation and

$$\Delta \log C_t + \frac{1}{4}\log(1+R_t) - \frac{4+1}{2}(\Delta \log C_t)^2 = e_t \tag{14}$$

for the second order approximation. Our instrument validity test is a t-test obtained from the regression of these constructed errors on our instruments. The null is that the errors are uncorrelated with the instruments: a significant t-statistic would suggests that instruments are not valid. We establish instrument relevance (instrument weakness) using the Cragg-Donald F statistic. The null hypothesis of this test is instruments are jointly weak.

The instruments we consider are the ones used extensively in the literature[9]: the lagged interest rate, lagged consumption growth and lagged income. We add lagged consumption growth squared to the instrument set for the second order approximation. Table 4 reports the mean t-statistics for each instrument (for instrument validity) and a Cragg-Donald F-statistic for the overall relevance to the instruments of the first order approximation.

All reported results are the mean values of the corresponding tests obtained through Monte-Carlo experiments. These experiments are performed using the simulated consumption paths of 1000 individuals (observed 40 periods) drawn from population of 10,000 individuals 1000 times with replacement. All individuals face the same interest rate realizations but individual consumption paths differ due to different income realizations. A simulated population of 10,000 individuals is generated after solving the life cycle model described in the Appendix A.

[9] We use the same instrument set as Attanisio and Low (2004).

Table 4 shows that in the patient versions of the models (1, 3, 5), where effective curvatures is low, there is no evidence that the instruments are invalid. However, when we examine the patient versions of the models (model 2 and 4), there is strong evidence that, lagged interest rate, the instruments are invalid. For example, increasing the discount factor in the Attanasio-Low model (moving from model 1 to model 2) results in higher effective curvature (from 0.03 to 0.48) and although the lagged interest rate remains valid, lagged consumption growth and lagged income become invalid instruments. The same pattern is observed also when we compare models 3 and 4. For instance, the t-statistic of lagged consumption growth in Model 3 moves from -0.8 to -18 in Model 4. This pattern changes dramatically for the Deaton models. Due to kink, the impatient version of this model has a very high effective curvature, but instrument validity is almost the same as the patient model (5).This is not surprising, since for estimating this model we remove the observations that are causing instrument validity problems (observations in which the Euler equation does not hold). As we will discuss in the next subsection, although there isn't an instrument problem, this procedure also results in low performance of the estimator. The last column in Table 4 presents the Craig-Donald F statistic for instrument relevance. As clear from the Table 4, instrument weakness is not an issue for the first order approximation. All the instruments are relevant. To summarize, we find that first order approximations do not face serious instrument problems in the low effective curvature environments.

The overall picture is quite different when we move to the second order approximation. An obvious justification for including higher order terms is to reduce the

correlation between the instruments and the composite errors. In addition, including variance of consumption growth allows the estimation of the relative prudence parameter and this in turn allows for a test of the iso-elastic sub-utility form (since in this form relative prudence and EIS are governed by a single parameter). One can in principle include as many higher order terms as possible. However inclusion of higher terms also necessitate of the use new instruments, and all possible instruments faces with augmented instrument problems. [10] This can be clearly seen in Table 5. While instruments are fairly valid, again especially models with low effective curvatures, they become irrelevant (weak). This is clear for models 1, 5 and more strongly for model 6. Moving from model 1 to model 2 (moving to higher effective curvature environment) makes the instruments jointly relevant but strongly invalid. The same trade off can also be observed in Carroll models (3 to 4).

While model 5 acts exactly like model 1, model 6 yields the most intriguing results. In this model the tradeoff between validity and relevance is the most clear. Although all the instruments seem valid, they have no predictive power. Why is this model so different? As we described, this is a model where the explicit borrowing constraint sometimes binds. This means that agents in their life-time operate (move back and forth) around the kink in the consumption function. Even though we remove the periods when they operate on the 45 degree line (periods in which agents don't carry any assets forward to the next period), curvature around the kink is still extremely deep and leads to high effective curvature. At the same time, with this selection process we

[10] Also including higher order terms of noisy variables is likely to cause significant small sample bias.

effectively remove the periods when lag values have strong predictive power for subsequent periods. This explains the extreme weakness (and validity) of the instruments in this model.

Table 5 establishes the relationship between the degree of effective curvature, instrument validity and instrument relevance for a given life cycle model. The overall picture shows a trade-off between instrument validity and relevance. For models with higher degree of effective curvature, instruments are more relevant but less valid. On the other hand for patient models, where the effective curvature is very low, instruments are valid but less relevant. In order to illustrate this point better consider models 1 and 5. In these environments agents are patient individuals who are likely to accumulate wealth. For these individuals consumption decision is largely based on financial wealth rather than their labor income. Therefore the second order term which is mostly related to idiosyncratic (labor income) uncertainty is not relevant for these agents. However, these individuals respond strongly to interest rate variations. As a result, while the instruments used in these models have power to identify the EIS, they are extremely weak with respect to prudence. The reason why irrelevance is not as strong for model 3 (the patient version of Carroll model) is the nature of the assumed background risk. The fact that individuals can face a zero labor income at any time in their life cycle heightens prudence considerably even if they are patient. Therefore higher order moments remain important.

Given these results, it would be informative to see how seriously the instrument problem manifests itself in actual estimates.

2.3 How serious is the bias?

In this subsection, we present our small sample results from the GMM and LIML[11] estimation of first and second order approximate Euler equations. As mentioned above, the estimates from these equations are either biased (if the instruments are invalid) or weakly identified, which is associated with small sample bias (if instrument relevance is the issue). Here, we investigate the magnitude of this bias in different simulated environments using Monte-Carlo experiments[12].

Table 6 reports results from the estimation of first order approximate Euler equations. Average estimates of the EIS and average standard errors of that parameter are reported. The third row of each model reports the percentage of times the true value of EIS (0.25) falls inside the 95% confidence interval. It is clear from the Table that there is almost no difference between GMM and LIML estimates. This is not surprising since instrument weakness is not issue in the first order approximation. Clearly, in all models the patient versions of the models are doing much better than their impatient counterparts (although mean estimates are close to true value in all models).

Table 7 presents estimation results for the second order approximation. Here, in addition to EIS we also estimate the prudence parameter (the true value is 2.5). As we established in the previous section, instrument weakness becomes a serious issue in this specification. Results for the EIS are quite similar to those from the first order approximation. Coefficients are significant and have negligible biases. Results for the

[11] We use LIML since it is known to perform better with weak instruments (see Murray 2006).
[12] Monte-Carlo experiments are performed using the simulated consumption paths of 1000 individuals (observed 40 periods) drawn from population of 10,000 individuals 1000 times with replacement.

prudence parameter are quiet different across models. For models 1 and 5 the estimates are not significantly different from zero. Statistically, this result is the artifact of the weak identification. Also it is not surprising to find that patient wealth accumulators do not respond to idiosyncratic uncertainty since these agents have enough wealth to perfectly smooth their consumption. For impatient models, the estimated prudence parameters are significant. However instrument invalidity prevents them from being accurately estimated. The only model that comes close to the true value of prudence is Carroll's impatient model (4). Given the strong invalidity of the instruments we repeat this particular experiment (model 4) by setting the coefficient of relative risk aversion to 2 while keeping all other parameters unchanged. This case naturally generates lower lifetime wealth. As reported in Table 8, the instrument invalidity problem becomes much more pronounced. Except for lagged income all other instruments are invalid. For example, here the t-statistic of the lag consumption growth is -11.05 as compared to -3.06 in the previous specification. Also, in this specification the lagged interest rate becomes an invalid instrument. Not surprisingly, the instrument relevance statistic more than doubled. As a result, even though for model 4 we came close to capturing the true value of the prudence parameter, this result is not robust to alternative CRRA values.

We also examine the effect of panel length on our results. Since, in the next section we present an empirical application with (at most) 14 years of data, we repeat all our experiments for panel length of 14 instead of 40. [13] Our results verify the importance of long panels. We find that as the sampling period shortens, both instrument invalidity

[13] Results for instrument test and first order estimation are not reported but they are available upon request

and relevance problems become more severe. What is striking is that estimated EISs, especially in the patient models, are again very close to the true values. For the second order approximation, the instrument relevance problem becomes more severe and plagues the estimations. In fact, only model 4 passes the relevance test, but again in this model instruments are invalid. The estimates are similar to 40 period experiments; the prudence parameter estimate close to the true value only for model 4 (see Table 9).

An important caveat applies here. The fact that we obtain simulation results that are favorable toward approximate Euler equation estimation (especially first order), even under short panel length, may be due to assuming away aggregate shocks to income. As Attanasio and Low (2004) stated, under the aggregate shocks one needs sufficiently long panels to make sure the effect of the shocks averages out over time.

Overall, Monte-Carlo results suggest that identification of EIS is possible with perhaps a small bias even when using the second order approximations. The bias in the estimated EIS is smaller for the patient models (where effective curvature is low). The identification of prudence on the other hand is not possible in most environments[14]. When patient models (1, 3 and 5) are used for prudence estimation, instruments are valid but not relevant. This leads to weak identification. When Carroll type buffer stock savers are sampled (model 4) the response to uncertainty becomes significant but instrument invalidity plagues the estimations. Surprisingly the estimated coefficient of prudence

[14] Of course, with the iso-elastic utility form, a good estimate of γ from the interest rate term also helps to identifiy the prudence (since it is equal to $\left(\frac{\gamma+1}{2}\right)$). However, here our point is, we are not able to get reliable estimate of γ from the the volatility of consumption growth.

does not exhibit any significant bias. When the Deaton type buffer stock savers are sampled (model 6), instrument irrelevance leads to a large bias in estimated prudence. In the next section we perform an empirical application using household survey data to see if real data support our simulation results.

3. An Empirical Application

In this section we estimate the first and the second order approximation to consumption Euler equation using the American Panel Study of Income Dynamics (PSID). The PSID contains annual information on food at home and food at restaurants. Despite its shortcomings (no expenditure variable other than food, large measurement error, lack of representativeness due to attrition) we chose to work with the PSID because it is the longest available panel survey on consumption and it has been used extensively for Euler equation estimation previously. Our sample covers the periods between 1974 and 1987. Although the actual panel length is much longer, some of the food variables are hard to interpret prior to 1974 and food related questions were suspended for several years after 1987.

We eliminate the original "poverty sample" from the data to make our sample more representative of the U.S. households[15]. Households headed by a student or someone older than 65 at any time of the sample period are excluded. Families that have consumption growth outside the 0.25-4 range in any period are eliminated. We include split-off families in our sample as long as they appear as a new family before 1974. We

[15] PSID contains two main samples: The Survey of Research Center (SRC) cross sectional sample and Survey of Economic Opportunity sample, the latter was upper truncated in order to study poverty issues.

do not let any split off's in our sample period since their initial years mostly cause noise in our data. Our sampling scheme is designed to pick out at least five consecutive periods in which food expenditure information was reported. Thus we have at least four consecutive years for each sampled households in which we can observe current and lagged consumption. In order to make our sample as similar as possible to the ones in simulations we impose a further restriction: To replicate the model 1, 3 and 5 we sample households that carry forward positive assets for several lagged periods. In particular, an observation in period t is included if at t-1, t-2 , t-3 and t-4 (4 consecutive lagged periods) a household did have savings equal or more than 2 months of her income[16].These constitute our "wealthy" households. Unfortunately, data do not allow us to form a buffer stock sample similar to models 2, 4 or 6. Our final unbalanced panel (maximum of 10 years) for wealthy sample has a total of 446 households and 2959 observations. As in the simulations, we assume that all households face a common real interest rate series calculated from the U.S. three-month treasury bill rate and the consumer price index. This amounts to using only the time variation in the intertemporal price.

3.1 Empirical Results

Table 10 reports our empirical results from PSID. Since the objective of this exercise is to test our findings in the survey data, we use the same instruments[17] and estimation methods as in the simulation exercises. We only add the first difference of family size in

[16] We use the same approach as in Zeldes (1989) to construct this threshold.
[17] In order to deal with measurement error, we used twice lagged of them instead of first lags.

to the right side of the first and second order specifications[18]. As it can be seen from the first column of the Table; the first order approximation on the wealthy sample yields a significant positive EIS (1.03 with standard error of 0.59). The second order approximation results are more in line with our simulation results. The EIS estimate is still positive significant (1.04 with standard error of 0.60). Turning to the prudence estimates, the estimates are negative and insignificant.

Overall our empirical findings underline the importance of sampling when using approximate Euler equations. Although instrument invalidity and irrelevance seem to be significant problems, especially with short panel length, we show that careful sampling (sampling wealthy households) can lead to better estimates of the EIS. On the contrary, we find that identifying the response to idiosyncratic uncertainty is not possible with available data. As shown in the simulations, wealthy households do not have a significant response to uncertainty, so identifying this parameter using this sample is doomed to failure.

3.2 A Discussion: What else might be going on?

The fact that it is more likely to extract a response to interest rates among wealthy households and extract a response to uncertainty among buffer stock savers can be thought of within a consumption floor framework. Suppose that households have a natural consumption floor below which they cannot survive. If we modify the sub-utility function accordingly we have the following exact Euler equation

[18] This is done since marginal utility of consumption of household clearly depends on family size.

$$\left(\frac{C_t-\theta}{C_{t-1}-\theta}\right)(1+R_t)\beta=\varepsilon_t \text{ with } E_{t-1}(\varepsilon_t)=1$$

where θ is the reference consumption level. The corresponding second order linear approximation will be

$$\Delta\log C_t-\frac{1}{\gamma}\left(\frac{C_t-\theta}{C_t}\right)\log(1+R_t)+\frac{\gamma+1}{2}\left(\frac{C_t}{C_t-\theta}\right)(\Delta\log C_t)^2=e_t$$

Now, suppose that C_t approaches θ. This would be the case for less wealthy households (buffer stock savers). In this situation the response to interest rate tends to disappear whereas the response to uncertainty magnifies. Conversely, if C_t diverges from θ (which is the case for wealthy households) the response to the interest rate becomes stronger and the response to uncertainty disappears. Note that, this possibility does not preclude the effects of approximation bias and invalidity/relevance of instruments. The idea is in a world where there exists a consumption floor, identification of the intertemporal allocation parameters suffers from an additional problem.

4. Conclusion

A large and growing empirical literature investigates the sensitivity of consumption (or, equivalently, saving) to interest rates (the elasticity of intertemporal substitution) and to uncertainty (prudence). Most empirical work estimates these effects by employing the optimality condition (Euler equation) derived from the dynamic optimization problem of a consumer. Often these non-linear equations are log-linearized (approximate Euler equations) so that linear instrumental variables methods can be used

to deal with measurement errors in consumption. However, it has recently been argued that the approximation bias induced by linearization may be worse than the problems that linearization is intended to solve. Consequently, the usefulness of the Euler equation approach has been called into question (eg., Carroll, 2001). We have explored this issue with a series of simulation studies. We showed that: (1) the problems associated with approximate Euler equation are strongly related to the assumed economic environment (in particular, the degree of approximation bias is related to the concavity of the underlying policy rules, which in turn follows from parameters of the economic environment); (2) there exist environments where approximate Euler equation can yield unbiased estimates (as well as environments in which the approximate Euler equation works quite badly); and, (3) careful sample selection (or conditioning) can lead to better estimates of intertemporal allocation parameters in a range of economic environments.

References

Alan, S and M. Browning, (2006), "Estimating Intertemporal Allocation Parameters using Simulated Expectation Errors," Economics Series Working Papers 284, University of Oxford, Department of Economics.

Altonji, J. G. and A. Siow, (1987), "Testing the Response of Consumption to Income Changes with Noisy Panel Data," Quarterly Journal of Economics, 102, 293-328.

Attanasio, O. P. , J. Banks, C. Meghir, G. Weber, (1999), "Humps and Bumps in Lifetime Consumption," Journal of Business and Economic Statistics, 17(1), 22-35.

Attanasio, O. P. and H. Low, (2004), "Estimating Euler Equations," Review of Economic Dynamics, 7(2), 405-435.

Attanasio, O. P. and G. Weber, (1993), "Consumption Growth, the Interest Rate, and Aggregation," Review of Economic Studies, 60, 631-49.

Bond, S. and C. Meghir, (1994), "Dynamic Investment Models and the Firm's Financial Policy," Review of Economic Studies, 61(2), 197-222.

Carroll, C., (2001), "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)," Advances in Macroeconomics, 1(1), 1003-1003.

Deaton, A., (1991), "Saving and Liquidity Constraints," Econometrica, 59(5), 1221-1248.

Dynan K. E., (1993), "How Prudent Are Consumers?" Journal of Political Economy, vol. 101, no:6, 1104-1113.

Dynan K. E, (2000), "Habit Formation in Consumer Preferences: Evidence from Panel Data," American Economic Review, vol. 90 (June 2000), 391-406

Hall, R. E., (1978), "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," Journal of Political Economy, December, vol. 86, 971-987.

Hansen, L .P. and K. J. Singleton, (1982), "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," Econometrica, 50(5), 1269-1286.

Ludvigson, S. and C. Paxson, (2001), "Approximation Bias in Linearized Euler Equations," Journal of Economics and Statistics, vol. 83, 2, 242-256.

Mehra, R. and Prescott, E. C., (1985), "The equity premium: A puzzle," Journal of Monetary Economics, 15(2), 145-161

Mulligan, C., (2004), "What Do Aggregate Consumption Euler Equations Say about the Capital-Income Tax Burden?" American Economic Review, 94(2), 166-170

Murray, M. P., (2006), "Avoiding Invalid Instruments and Coping with Weak Instruments," Journal of Economic Perspectives, 20(4), 111-132.

Runkle, D. E., (1991), "Liquidity Constraints and The Permanent Income Hypothesis," Journal of Monetary Economics, 27, 73-98.

Stock, J. and Yogo, M., (2000), "Testing for Weak Instruments in Linear IV Regression," NBER Technical Working Papers 0284.

Yogo, M., (2004), "Estimating the Elasticity of Intertemporal Substitution When Instruments Are Weak," The Review of Economics and Statistics, 86(3), 797-810

Zeldes, S. P., (1989), "Consumption and Liquidity Constraints: An Empirical Investigation," Journal of Political Economy, 97(2), 305-46

Tables

Table 1: Parameter Values

Parameter	Value
Coefficient of Risk Aversion (γ)	4
Discount Rate $\left(\delta=(\frac{1}{\beta})-1\right)$	0.03 and 0.07
Standard Deviation of Permanent Income Shocks (σ_n)	0.05
Standard Deviation of Transitory Income Shocks (σ_u)	0.1
Unconditional Mean of Interest Rate Process (μ)	0.03
AR(1) Coefficient of Interest Rate Process (ρ)	0.6
Standard Deviation of Interest Rate Process (σ_ε)	0.025
Probability of Zero Income (Models 3 and 4)	0.01

Table 2: Models

Model	Impatient	Patient	Borrowing Constraint
1	-	Yes	No
2	Yes	-	No
3	-	Yes	Implicit
4	Yes	-	Implicit
5	-	Yes	Explicit
6	Yes	-	Explicit

Table 3: Moments of Expectational Errors

Model	Effective Curvature	Var	Skw	Krt
1	0.0302	0.03	0.522	3.55
2	0.4801	0.061	1.01	6.11
3	0.0852	0.039	14.56	2233
4	1.673	0.237	55.38	5788
5	0.0365	0.03	0.523	3.55
6	2.667	0.047	0.914	5.8

Note to Table 3:

- Reported statistics are the mean values of 1000 Monte Carlo

Table 4: Instrument Validity and Relevance Results, First Order Approximation

$$\Delta \log C_{t+1}^h = \alpha - \frac{1}{\gamma}\log(1+R_{t+1}) + e_{t+1}^h$$

	Instrument Validity				Instrument Relevance
	$\alpha + e_{t+1}^h = \Delta \log C_{t+1}^h + \frac{1}{\gamma}\log(1+R_{t+1})$				F
	Mean t stat				[10% , 90%]
	$\log(1+R_t)$	$\Delta \log C_t^h$	$\log y_t^h$	Mean R^2	
Model 1	-0.916	-0.459	0.655	0.00004	7438
(0.0302)					[7431 , 7447]
Model 2	-1.042	-6.746	-4.774	0.00175	7457
(0.4801)					[7442 , 7478]
Model 3	0.448	-0.754	0.619	0.00004	7362
(0.0852)					[7321 , 7402]
Model 4	-0.928	-18.254	1.58	0.00914	7362
(1.673)					[7320 , 7402]
Model 5	-0.957	-0.476	0.673	0.00004	7438
(0.0365)					[7431 , 7447]
Model 6	-1.614	-1.424	-1.074	0.00015	6962
(2.667)					[6853 , 7076]

Notes to Table 4:

1. Instruments used are lag interest rate , lag consumption growth and lag income.
2. Mean values of F and t statistics are reported. F values at 10 and 90 percent are in parentheses

Continues on the next page

3. Instrument Relevance column reports the first stage F statistic from an endogenous variable Interest Rate ($\log(1+R_{t+1})$), on to instruments lagged interest rate($\log(1+R_t)$),lagged consumption growth ($\Delta \log C_t^h$) and lagged income ($\log y_t^h$).
4. Stock and Yogo (2002) - test for weak instrument H_0=bias of two stage estimation relative to OLS is greater than 10% Critical Value at 5% significance level when the number of instruments is 3 for Linear GMM = 9.08, for LIML=6.46.

Table 5: Instrument Validity and Relevance Results, Second Order Approximation

$$\Delta \log C_{t+1}^h = \alpha - \frac{1}{\gamma}\log(1+R_{t+1}) + \frac{\gamma+1}{2}(\Delta \log C_{t+1}^h)^2 + e_{t+1}^h$$

	Instrument Validity					**Instrument Relevance**
	$\alpha + e_{t+1}^h = \Delta \log C_{t+1}^h + \frac{1}{\gamma}\log(1+R_{t+1}) - \frac{\gamma+1}{2}(\Delta \log C_{t+1}^h)^2$					Cragg-Donald F statistic
	Mean t stat					[10% , 90%]
	$\log(1+R_t)$	$\Delta \log C_t^h$	$\log y_t^h$	$(\Delta \log C_t^h)^2$	Mean R^2	
Model 1	-2.002	-0.537	-0.149	-0.41	0.000127	7.49
(0.0302)						[1.774 , 15.329]
Model 2	-1.817	-4.209	-2.41	2.917	0.000929	92.843
(0.4801)						[50.507 , 163.971]
Model 3	-0.417	0.276	-0.208	-4.316	0.004666	409.125
(0.0852)						[3.289 , 5308]
Model 4	0.001	-3.062	-0.103	-4.811	0.00277	2008.187
(1.673)						[555.01 , 5400]
Model 5	-2.024	-0.544	-0.105	-0.521	0.000133	7.457
(0.0365)						[1.999 , 15.525]
Model 6	-0.971	-0.917	-0.938	0.34	0.000075	2.636
(2.667)						[.156 , 9.437]

Notes to Table 5:

1. Instruments used are lag interest rate, lag consumption growth, lag income and lag consumption growth square.
2. Mean values of F and t statistics are reported. CDF values at 10 and 90 percent are in parentheses

Continues on the next page

3. Instrument Relevance column reports the CDF from the first stage of IV. Endogenous variables Interest Rate ($\log(1+R_{t+1})$) and lagged consumption growth square $\left(\Delta \log C_{t+1}^{h}\right)^2$, are regressed on to instruments lagged interest rate($\log(1+R_t)$),lagged consumption growth ($\Delta \log C_t^h$),lagged consumption growth square $\left(\Delta \log C_t^h\right)^2$ and lagged income ($\log y_t^h$).
4. Stock and Yogo (2002) - test for weak instrument H_0=bias of two stage estimation relative to OL S is greater than 10% Critical Value at 5% significance level when the number of instruments is 4 for Linear GMM = 7.56, for LIML =4.72

Table 6: Monte Carlo Results for the First Order Approximation

Estimates of the EIS using First Order Approximation $\Delta \log C_{t+1}^h = \alpha - \frac{1}{\gamma}\log(1+R_{t+1}) + e_{t+1}^h$		
True Value of EIS $\left(\frac{1}{\gamma}\right)$ is 0.25		
	Linear GMM	**LIML**
Model 1	.239	.2389
(0.0302)	(.012)	(.0119)
	83.1	83.4
Model 2	.2341	.2389
(0.4801)	(.0169)	(.0167)
	85.2	90.8
Model 3	.2545	.2546
(0.0852)	(.0117)	(.0117)
	93.2	93
Model 4	.2237	.2203
(1.673)	(.0151)	(.0148)
	59.9	48.5
Model 5	.2385	.2384
(0.0365)	(.012)	(.0119)
	84.7	84.8
Model 6	.2246	.2247
(2.667)	(.0156)	(.0153)
	64.9	63.8

Note to Table 6 :

- Mean values of estimates are reported , mean standard errors are in the parentheses. Third row of each model reports the percentage of times real value of parameter is inside the 95% confidence interval.

Table 7: Monte Carlo Results for the Second Order Approximation

Estimates of the EIS and Prudence Using Second Order Approximation

$$\Delta \log C_{t+1}^h = \alpha - \frac{1}{\gamma}\log(1+R_{t+1}) + \frac{\gamma+1}{2}(\Delta \log C_{t+1}^h)^2 + e_{t+1}^h$$

Real Value	**EIS** $\left(\frac{1}{\gamma}\right)=0.25$		**Prudence** $\left(\frac{\gamma+1}{2}\right)=2.5$	
	Linear GMM	LIML	Linear GMM	LIML
Model 1	0.2295	0.23	1.705	1.532
(0.0302)	-0.0207	-0.023	-3.069	-3.558
	83.6	83.8	95.7	95.3
Model 2	0.2051	0.205	5.928	5.969
(0.4801)	-0.0195	-0.0196	-0.7667	-0.6943
	36	36.8	0.2	0.2
Model 3	0.2459	0.2458	1.973	2.021
(0.0852)	-0.0147	-0.0145	-1.705	-1.669
	93.1	92.7	79.6	76
Model 4	0.2463	0.2446	2.526	2.341
(1.673)	-0.0198	-0.0185	-0.3673	-0.0708
	92.3	90.1	78.9	19
Model 5	0.2296	0.2303	1.646	1.505
(0.0365)	-0.0207	-0.0224	-3.095	-3.448
	82.1	82.9	96.2	95.3
Model 6	0.2489	0.2593	6.49	9.856
(2.667)	-0.0326	-0.8925	-6.397	-167
	93.8	96.6	99.2	98.6

Note to Table 7 :

- Mean values of estimates are reported , mean standard errors are in the parentheses. Third row of each model reports the percentage of times real value of parameter is inside the 95% confidence interval.

Table 8: Model 4 Revisited
(Second Order Approximation)

	Instrument Validity					**Instrument Relevance**
	Mean t-stat					Cragg-Donald F statistic
	$\log(1+R_t)$	$\Delta \log C_t^h$	$\log y_t^h$	$\left(\Delta \log C_t^h\right)^2$	Mean R^2	[10% , 90%]
Model 4 ($\gamma = 4$)	0.001	-3.062	-0.103	-4.811	0.00277	2008.187 [555 , 5400]
Model 4 ($\gamma = 2$)	2.7	-11.05	-1.01	-20.24	0.01343	4525 [2763 , 5537]

	Estimation			
	EIS		**Prudence**	
	Linear GMM	LIML	Linear GMM	LIML
Model 4 ($\gamma = 4$)	$\left(\frac{1}{\gamma}\right) = 0.25$		$\left(\frac{\gamma+1}{2}\right) = 2.5$	
	0.246	0.245	2.53	2.34
	-0.02	-0.02	-0.367	-0.071
	92.3	90.1	78.9	19
Model 4 ($\gamma = 2$)	$\left(\frac{1}{\gamma}\right) = 0.5$		$\left(\frac{\gamma+1}{2}\right) = 1.5$	
	0.453	0.437	1.35	1.15
	-0.043	-0.04	-0.075	-0.018
	25.6	27.6	45.6	0.2

Table 9: When Panel length is 14
(Second Order Approximation)

	EIS		**Prudence**	
Real Value	$\left(\frac{1}{\gamma}\right)=0.25$		$\left(\frac{\gamma+1}{2}\right)=2.5$	
	Linear GMM	LIML	Linear GMM	LIML
Model 1	.2791	.2697	-3.292	-5.678
(0.0302)	(.0468)	(.8903)	(7.627)	(264)
	87.4	89.5	96.1	96.2
Model 2	.4021	.4029	4.428	4.648
(0.4801)	(.053)	(.0543)	(1.194)	(1.112)
	17.8	19.2	64.6	51.6
Model 3	.272	.272	-3.195	-3.008
(0.0852)	(.0417)	(1.827)	(5.274)	(649)
	89	87.3	91.1	90.1
Model 4	.2753	.2762	2.71	2.62
(1.673)	(.0492)	(.0493)	(.493)	(.1551)
	90.3	90.7	70.3	21.9
Model 5	.2772	.2671	-3.644	-7.118
(0.0365)	(.0464)	(.2361)	(7.329)	(54.677)
	87.3	89	95.3	95.7
Model 6	.3324	.3224	1.635	1.754
(2.667)	(.0671)	(5.683)	(9.82)	(1089)
	72.8	77.9	97	97.3

Note to Table 9 :

- Mean values of estimates are reported , mean standard errors are in the parentheses. Third row of each model reports the percentage of times real value of parameter is inside the 95% confidence interval.

Table 10: Estimates from the PSID
(Rich Sample)

	First Order		Second Order	
	Linear GMM	LIML	Linear GMM	LIML
$\log(1+R_{t+1})$	.9978	1.033	1.041	1.051
	(.6006)	(.595)	(.6045)	(.5967)
$(\Delta \log C^h_{t+1})^2$			-.1534	-.1547
			(.2248)	(.1878)
Δ in Fam. Size	.1101	.1104	.1112	.1116
	(.0128)	(.0122)	(.0129)	(.0123)
Constant	-.0335	-.0361	-.0186	-.0187
	(.0193)	(.0192)	(.0294)	(.0283)
Sample Size	2959		2959	
Cragg-Donald	225		18	

Notes to Table 10:

1. ***First Order Instruments are*** $\Delta \log C^h_{t-1}$, $\log(1+R_{t-1})$, $\log(y^h_{t-1})$
2. ***Second Order Instruments are*** $\Delta \log C^h_{t-1}$, $\log(1+R_{t-1})$, $\log(y^h_{t-1})$, $(\Delta \log C^h_{t-1})^2$

Figure

Figure 1:Consumption Functions and Distribution of Cash-on-Hand

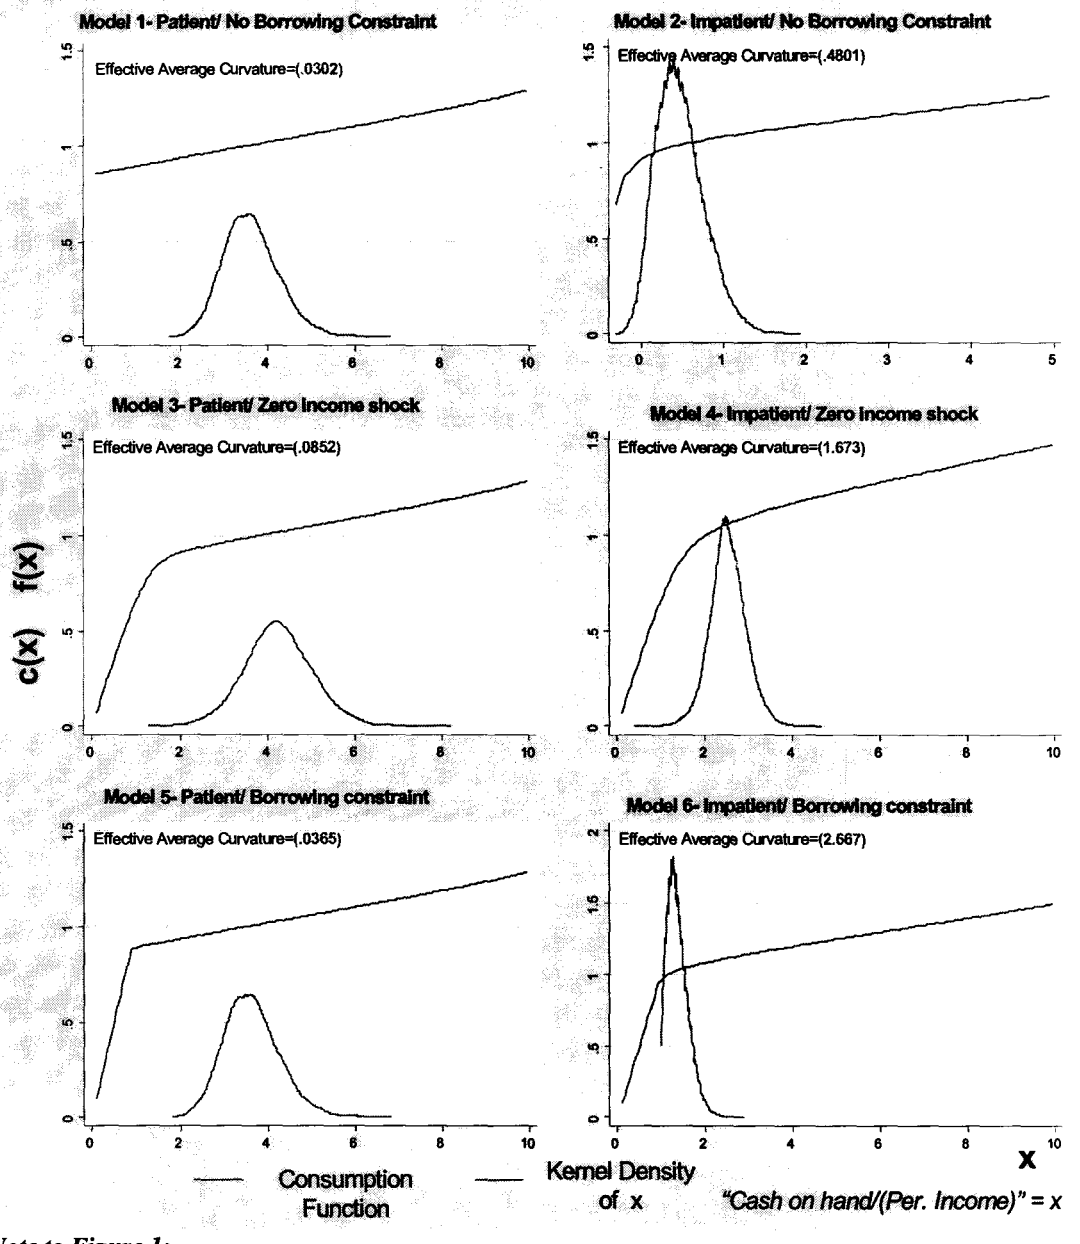

Note to Figure 1:

- Consumption Function and distribution of normalized cash-on-hand at the age of 40. Non-parametrically calculated effective curvature measures are in parentheses. X axis is cash-on-hand to permanent income ratio

Appendix

The Life Cycle Model

We assume that the utility function is intertemporally additive and the sub-utilities are iso-elastic. The problem of the generic household is:

$$\max_{c_t} E_t \sum_{t=0}^{T} \beta^t u(C_t)$$

$$s.t \quad X_{t+1} = (1 + R_{t+1})(X_t - C_t) + Y_{t+1}$$

where C_t is nondurable consumption in period t, X_t is cash-on-hand (total financial and nonfinancial wealth) and Y_t is current labor income. We assume that durable consumption and leisure are separable from the nondurable consumption. The income process is assumed as follows:

$$Y_{t+1} = P_{t+1} U_{t+1}$$

$$P_{t+1} = G P_t N_{t+1}$$

where G is predictable permanent income growth, P_t is permanent income which is subject to log-normally distributed shocks N_t with mean unity and variance $\left(e^{\sigma_n^2} - 1\right)$, current income Y_t equals permanent income multiplied by a transitory shock, U_{t+1}, which is distributed lognormally with mean one and variance $\left(e^{\sigma_u^2} - 1\right)$. The interest rate series is assumed to be generated by a stationary first order autoregressive process with long-

run mean μ and autoregressive coefficient ρ. Interest rates shocks ε_{t+1} are assumed to be white noise with variance σ_ε^2. The process is

$$R_{t+1} = (1-\rho)\mu + \rho R_t + \varepsilon_{t+1}$$

The intertemporal model described above does not have an analytical solution due to the assumed income uncertainty. Therefore we utilized the standard numerical dynamic programming methods to obtain a solution. Since the utility function is additive over the life cycle we solved the model recursively starting from the last period of life. We assume away any bequest motive so that consumption in period T is:

$$c_T(x_T) = x_T$$

The problem is solved via policy function iteration using the terminal value condition. Having a nonstationary income process makes the problem harder to solve since the range of possible income values is too large. Instead, we redefine all relevant variables in terms of their ratios to permanent income and solve for the consumption to income ratio. By doing this we reduced the number of state variables to two, namely the cash on hand to income ratio and the interest rate. Moreover, we obtain an iid income process which can be approximated by standard Quadrature methods. Given the redefinition of the variables, the Euler equation can be written as

$$\theta_t(\omega_t, R_t)^{-\gamma} - \frac{1}{1+\delta} E_t\left[(1+R_{t+1})\theta_{t+1}(\omega_{t+1}, R_{t+1})^{-\gamma} n_{t+1}^{-\gamma}\right] = 0$$

where $\theta_t = \frac{C_t}{P_t}$, $\omega_t = \frac{X_t}{P_t}$. At the terminal date T, consumption to income ratio is a function of only the cash on hand to income ratio and since the bequest motive is assumed away it follows that $\theta_T = \omega_T$

For the income process, we use 10 point Gaussian Quadrature and we approximate the interest rate process by forming a 10 point first order discrete Markov process. We use a cubic spline to approximate the consumption function at each iteration. The agent is allowed to borrow the amount he can pay back with certainty. In practice this constraint will never bind because the functional form of the utility function implies that zero consumption results in infinite marginal utility. In models where we do not assume an explicit borrowing limit (model 1, 2, 3 and 4), the consumption functions are continuously differentiable. In fact, in our case where agents have iso-elastic preferences and income uncertainty, consumption functions are strictly concave.

In order to solve the problem, we define an exogenous grid for the cash on hand to income ratio: $\{x_j\}_{j=1}^{J}$. It is important to adjust the grid as the solution goes back in time. The algorithm finds the consumption level that makes the standard Euler equation hold for each value of x and r. We made the grid for x finer at lower levels in order to capture the curvature of the consumption function. After solving for the consumption function of a generic household for 60 periods, we simulate consumption paths for 10,000 ex-ante identical households facing the same interest rate realizations. We use only the middle periods for the estimations. Table 1 presents the assumed parameter values for our experiments.

VI

Increasing the Saving Rates of Low Income Families

1. Introduction

Is it really feasible to get poorer families to save? And if so how? This essay sketches academic research which is relevant to these questions. In the current Canadian policy context, the target group is families with children, in the bottom decile of the income distribution – below about $25,000 per annum. Although most academic research is not tightly focused on this particular group, it is possible to draw some useful lessons from the broader literature, and to make informed guesses about how the conclusions of this literature might need to be modified for this particular group.

Empirical studies using micro data suggest that low income families in Canada save very little. There are two ways to assess household savings. One is to examine the flow in a given period: either income minus expenditure or the change in net wealth. This is most often expressed as a rate (dividing by income). Note that income minus expenditure will differ from the change in net wealth over a period because of capital gains (or losses). Income minus expenditure is "active saving", while changes in net wealth will reflect both "active saving" and capital gains or losses ("passive saving"). The alternative to a flow measure is to look at wealth levels (the stock). Net wealth reflects both past active saving (and dis-saving) and past returns (capital gains and losses). Thus the stock summarizes information about saving behavior over a longer period (as well as about returns).

Using data on (active) saving flows constructed from the Canadian Family Expenditure Survey, Burbidge and Davies (1994) show that at all ages, household saving rates are negative in the lowest current income quintile (that is, the bottom 20%). Alan, Atalay and Crossley (2006) show that the saving rates of lowest income quintile of Canadian households remain very low (essentially zero) when households are classified by a measure of permanent (rather than current) income. This is important, for reasons that will be elaborated below. Turning to the stock of wealth, data from the 2005 Survey of Financial Security reveal that Canadian households in the bottom quintile of the (current) income distribution had median net wealth of just $1000, with a mean of -$2400 (StatCan, 2005). Thus by both stock and flow measures, low income families in Canada appear to save very little.

In thinking about whether it might be possible to raise the saving rates of low income families, it is useful to organize our discussion around three questions:

1. What does the traditional economic framework for thinking about savings – the life-cycle/permanent income hypothesis (LC/PIH) – say about this?
2. What about the emerging literatures on "choice architecture" and financial literacy?
3. Is there evidence that we can change preferences?

The LC/PIH framework assumes that preferences are fixed, and that individuals are forward looking and pretty good at optimizing. The literature on "choice architecture" takes preferences as fixed, but maintains that optimizing is hard. People use 'rules of thumb' and other heuristics, and they are sometimes 'lazy' or simply time-constrained, or

just forgetful. Consequently, their decisions are affected by the framing of choices, the setting of default options, and so on. Finally, some researchers and policy makers believe (or hope) that preferences are not fixed, and that by some intervention it might be possible to inculcate a 'taste for saving' or a saving 'habit'.

The first of these questions is addressed in the next Section, on the LC/PIH framework. The second and third questions are taken up in Section 3, on the "New Economics of Saving". Section 4 concludes with a summary of our main points.

2. The Life-Cycle/Permanent Income Hypothesis Framework

2.1 Description

The LC/PIH is the standard economic framework for thinking about saving decisions. It originated with the seminal work of Friedman (1957) and Brumberg-Modigliani (1954) (See also Modigliani, 1976), and has been substantially developed since. We consider the core of the LC/PIH to be the following assumptions:

a. Households have fixed, inter-temporally consistent preferences.

b. Households are forward-looking and have rational expectations in the sense that their forecasts of the future do not contain systematic errors (of course they make errors, but they do not make the same error repeatedly).

c. Households optimize. They do the best they can with the information available to them.

d. In any period, the marginal value of additional expenditure is declining. All else equal, an extra dollar of spending is more valued in a period which a household has consumed very little, than in a period in which the household has consumed a great deal.

Technically, assumption (a) will mean exponential discounting, and rule out models with hyperbolic discounting, which exhibit time-inconsistencies. Less technically, this means that when a household thinks about two dates in the future, A and B, the relative value it puts on some extra spending at those two dates is the same today as it will be when it arrives at date A. (For an introduction to hyperbolic discounting, see Angeletos et al., 2001)

Assumption (d) leads to the implication that households should 'smooth consumption', although this implication must be treated with care. What is actually implied by assumption (d) (in conjunction with the other assumptions) is that the household should keep the properly discounted value of the marginal utility of expenditure constant. This will only correspond to smoothing 'consumption' under additional assumptions, about preferences and whether the household faces uncertainty. This illustrates the point that a variety of auxiliary assumptions are always made to flesh out a model within the LC/PIH framework, for the purposes of analysis, simulation or estimation. However, the above list is all that we would consider central.

2.2 Modern Development

In the half-century since Friedman and Modigliani, economists have developed the LC/PIH substantially. The most important modern development is the formal

incorporation of uncertainty (Deaton, 1991, Carroll, 1997). This turns out to be very important. In particular, the idea (in Modigliani, for example) that the life-cycle path of consumption should be independent of the life-cycle path of income (except for satisfying the lifetime budget constraint) turns out to be a special result, essentially requiring either that households face no uncertainty (about future incomes, needs, or rates of return) or have very particular preferences (a specific utility function.) With more general preferences, the effect of uncertainty can be very much like liquidity constraints, causing consumption to 'track' income, especially early in life.

Some readers may have been surprised that our formulation of the LC/PIH did not contain an assumption of 'no liquidity constraints'. Because uncertainty (uninsurable risk) can induce similar behavior, it seems to us more sensible to think of the LC/PIH as a general proposition about preferences and behavior, without specifying the details of the economic environment. The behavior of LC/PIH savers will depend on whether they can borrow, what risks they can insure, and so on.

2.3 General Implications

What then are the implications of the LC/PIH for saving behavior? The LC/PIH implies that households will save when needs are low, income is high, and the rate of return is high. Conversely, they will tend to dissave when needs are high, income is low or the rate of return is low. For example, we would expect households to dissave when they are young (when needs are high because of children, and the acquisition of durables) and in retirement (when incomes are lower) and to save in the latter part of the working life (when incomes are typically higher, and needs are perhaps lower).

As with most of economics, the LC/PIH can be given a 'cost-benefit' interpretation, and this may be useful for thinking about policy. Households will save (forgo current consumption) up to the point where the cost of the last (marginal) dollar set aside equals the marginal benefit of saving that dollar. When current needs are high, or current income low, the 'cost' of setting aside a dollar for the future is high. If future income is expected to be high, or future needs low, then the marginal benefit of saving is low. If the available rate of return on saving is low, the marginal benefit of saving is low. Any of these three factors should be expected to lead to the lower saving.

2.4 Validity

Is the LC/PIH a theoretical framework with any validity? The first thing to note is that it is very easy to reject very simple, old-fashioned version of the LC/PIH with data. But this is not very interesting. Such models are straw men. It turns out to be rather more difficult to reject modern, sophisticated models in the LC/PIH framework. Browning and Crossley (2001) provide a detailed survey; a brief outline is as follows:

i. Attanasio, Banks, Meghir and Weber (1999) and Gourinchas and Parker (2002) have shown that a modern version of the LC/PIH (in particular, one with impatient households that face income uncertainty) replicates the broad pattern of life-time consumption (the well known 'hump') very well.

ii. An influential paper by Banks, Blundell and Tanner (1998) seemed to demonstrate a systematic failure to 'smooth consumption' across the retirement threshold (the so-called 'retirement-consumption puzzle'). However, more recent research suggests that households anticipate some fall

in *expenditures* with retirement (see Hurd and Rohwedder, 2006) and that this fall does not represent a fall in consumption or living standards, perhaps because they have more time to engage home production, or shopping for better prices, so that the same the same consumption can be obtained with less market expenditure (for some Canadian evidence on this point, see Brzozowski and Lu, 2006). Thus many economists no longer consider this drop to be evidence against the LC/PIH.

iii. Another class of tests examines whether households respond to anticipated income changes. The idea is that income changes which are known in advance should already have been incorporated into spending plans, and hence lead to no further change in spending when they occur. It is important to bear in mind that a failure to smooth expenditure is not necessarily a failure to smooth consumption (or marginal utility). For example, the fact that households tend to spend more around pay dates does not necessarily imply that they have declining consumption through the month: many goods are storable over this time horizon. Nevertheless, a reasonable reading of the literature here would be that the evidence is mixed: some researchers find responses and others do not. Browning and Crossley (2001) provide a reconciliation of this literature. They show that households deviate from the LC/PIH only when there is very little cost to doing so. When it matters, they get it right.

iv. Critics of the LC/PIH often point to age-wealth profiles as evidence against this framework. First, wealth profiles do not seem to fall sufficiently in old age, when the LC/PIH predicts dissaving. A second, related point is that households do not appear to hold enough wealth, particularly as they approach retirement. With respect to the first point it is important to note that typical survey data records dissaving from public and private pensions as 'income'. Thus dissaving is understated and because income (y) is overstated, saving (y-c) is inflated. The Canadian Survey of Financial Security is an unusual survey in that it contains data on the value of pension rights. Milligan (2005) shows that when this item is incorporated into a measure of total wealth, the age profile of wealth looks much more like what a LC/PIH model would predict. With regard to the second point, there is considerable controversy. The problem is that in order to go from the rather general LC/PIH (as outlined above) to a fully specified model which can be used to calculate how much wealth a given household should optimally hold at retirement, a great many auxiliary assumptions must be made. Different researchers make different assumptions (about how needs vary with household size, about the role of housing wealth, about the likely future path of medical or nursing home costs) and arrive at quite different conclusions. Contrast for example Bernheim et al (2000, 2001) with Engen et al. (1999) and Scholz et al. (2006). Skinner (2007) provides a very accessible discussion of the sensitivity of saving "adequacy" calculations to alternative assumptions.

A summary of these first four points is that it is very hard to reject to the LC/PIH with data on consumption or wealth.

v. On the other hand, the LC/PIH may have more difficulty explaining household portfolios. For example, many households appear to hold both liquid assets (earning little return) and high interest debt. This is difficult to reconcile in the LC/PIH framework. Angeletos et al. (2001) argue that this is an area where models with hyperbolic discounting (time-inconsistent preferences) are better able to fit the data.

vi. Finally, there is a growing list of "behavioural" tests that show individuals and households do make systematic errors of predictable kinds (for example, they are susceptible to manipulation though the framing of choices). That this is true is beyond dispute, and in this sense the LC/PIH is not entirely correct. However, the interesting questions are how costly these mistakes are, and what (if anything) policy makers should do about them. We return to these questions below.

2.5 Implications for Saving by the Poor - Interpreting Saving Data

Whether or not the LC/PIH is completely true, it offers important insights for interpreting data on savings. In fact, the LC/PIH was developed as a way of understanding data on income, saving and consumption, and in particular for resolving the puzzle of why saving rates rise with income in cross-sectional data, but average saving rates were fairly stable over time (as average incomes rose). One of the original and key insights of the LC/PIH is that a cross-sectional correlation between saving rates

and current income will be observed even if households that are "rich" and "poor" in a long run sense have similar saving rates. First, any measurement error in income will also increase observed saving if saving is measured as income minus consumption. Thus measurement error in income will lead to a completely spurious relationship between income and saving rates (or equivalently, between incomes and average propensities to consume, as the later is simply one minus the saving rate.) Second, if households experience transitory income fluctuations, the LC/PIH implies that households will save when their incomes are temporarily high, and dissave in periods when their incomes are temporarily low. Thus, if we have two households with the same average income, but with dis-synchronous income fluctuations, saving will be correlated with current income at any point in time. The implication is that if we wish to compare the savings rates of rich and poor households, it is important to try to group households by some proxy for "permanent income" – in which both measurement error and transitory income fluctuations are averaged out. Taking an average of incomes over many years immediately suggests itself, but while panel (or longitudinal) data on income are common, panel data on savings or consumption are very rare (moreover, such an average is only a good measure of permanent income if incomes are mean-reverting). Since we typically observe savings (and consumption) in only in cross-sectional surveys, a one period proxy for 'permanent income' is required. Figure 1, which is adapted from Alan, Atalay and Crossley (2006), is one attempt to do this. The underlying data are from Statistics Canada's Family Expenditure Survey for 1996. The saving rate, on the vertical axis, is a measure of the change in wealth, divided by annual income. Income is measured

on the horizontal axis. Households are divided into income quintiles (conditional on age) and each line traces out the median savings rate of the five income quintiles, for a different measure of income. The steepest line, which is blue, traces out the median savings rates of the quintiles of current income. Each higher quintile has a higher saving rate – savings rates rise rapidly with current income. All the other lines trace out the median savings rates of the permanent income quintiles. Permanent income is effectively imputed from household spending (formally, current income is used as an error-ridden measure of permanent income, which is then 'instrumented' with components of consumption.) The motivation for this procedure is that, if the LC/PIH is broadly true, then households should consume in proportion to their permanent income. Households with higher consumption probably have higher lifetime income. The instrumental variables procedure also deals with measurement error in income. These lines demonstrate two things. First, the median savings rates of the second, third, fourth and fifth income quintiles are very similar. This implies that much of the correlation between current incomes and savings rates has to do with measurement error and temporary income fluctuations. Secondly, and most importantly, this analysis suggests that those in the first quintile of permanent income (that is, those that really are poor in a long run sense) have very low savings rates.

Given this, the LC/PIH also provides a framework for thinking about why the poor save less (proportionally) than more affluent households. The LC/PIH says that saving behavior is determined by the timing of income and needs, and by rates of return. It therefore suggests that candidate explanations are that the poor face different age profiles

of income or needs, or that they face different returns. If the poor have low incomes in working life and access to means tested income guarantees in retirement, it may not be optimal for them to save for two reasons. First, their income profile is quite flat right into retirement (in other words, the replacement rate provided by the public components of the retirement income system is close to one – as documented by LaRochelle-Côté et al., 2008.) This means that it does not make sense to transfer income from early in life to latter in life. Second, the means testing of the retirement income guarantees may mean that the poor face very low effective rates of return on saving: any saving is 'taxed' away by the means testing in retirement. Shillington (2003) has suggested that this is an important issue in Canada, and Huggett and Ventura (2000) have explored this explanation for low savings rates among the poor in the U.S. Means testing in income support programs that operate in during the working life can also act to inhibit saving among the poor. This point has been made for the U.S. case by Hubbard, Skinner and Zeldes (1995). Note that both the Huggett-Ventura and Hubbard-Skinner-Zeldes papers are quite technical. They report the numerical solution (and then simulation) of a fully specified life-cycle model with uncertainty, in which detailed features of the U.S. social insurance programs have been programmed into the model. To the best of our knowledge, a similar exercise has not been conducted for Canada.

2.6 Implications for Saving by the Poor - Policy Options

Finally, the LC/PIH provides a guide to possible policy options. Perhaps the most important thing that the LC/PIH tells us is that it can be very harmful to compel people to save when the cost of saving is high or the benefit is low. Moreover, if the cost of saving

is high, or the benefit low, we should expect households to resist any attempt to compel them to save, by adjusting on unregulated margins (by running up more debt to offset the increased saving, for example.) Empirical analysis of compulsory savings regimes (such as the compulsory retirement savings programs in Australia and Chile) is an area of current research, but if the LC/PIH is broadly correct, then compulsion is likely to be either ineffective, or (if people do not have other margins to adjust) harmful.

Setting aside compulsion, the LCH/PIH framework says that any policy to increase saving must raise the benefit of saving or lower the cost. This means raising current income or lowering current needs or raising return on savings. The last is really the feasible policy lever, and governments around the world have attempted to use this lever repeatedly. Rates of return can be altered by, for example, tax-favoring savings. It is perhaps worth noting that with progressive income taxation, tax-favoring saving tends to raise the after-tax rate of return faced by richer households by more than it raises the after-tax rate of return of poorer households (a point made, for example, in Mills et al., 2006 and Duflo et al., 2006). This is one reason that recent policy initiatives have focused on matching contributions for targeted program participants rather than general tax favoring (these matched saving programs will be discussed further below, in the section on changing preferences.) Nevertheless, interest in tax-favoring and other manipulations of rates of return to increase saving has led to two large, and surprisingly inconclusive, literatures. One attempts to empirically estimate the key behavioral parameter: the "elasticity of inter-temporal substitution" which measures how sensitive consumption (and therefore saving) is to variations in the interest rate. The other goes directly at the

policy question and asks whether tax-favoured savings accounts have led to real increases in net saving.

With respect to the first literature, the issue is how to estimate the response of consumption (and hence saving) to interest rates when household level panel data on consumption are known to be very noisy, and we have little exogenous variation in the interest rate. Indeed, there is considerable disagreement about whether, given the measurement error in consumption data, the standard "Log-linearized Euler Equation" estimation approach is useful at all (contrast for example, Carroll, 2001 with Attanasio and Low, 2004). Among those who believe that the elasticity of intertemporal substitution can be estimated in this way, a reasonable range for the estimate is thought to be between 0.5 and 1. This is not terribly precise.

With respect to the second literature, the central issue is whether saving in tax-favoured vehicles represents net new saving, or simply the re-allocation of savings to these vehicles. Here again there is little consensus. For example, Poterba, Venti and Wise (2006) argue that saving in tax favored accounts in the U.S. is largely new saving while Engen, Gale and Scholz (1996) conclude the opposite. It is striking that these authors are employing the same data (though of course different estimation strategies). A good recent survey of both these literatures is Attanasio and Wakefield (2008).

2.7 Canadian Evidence on the Effectiveness of Savings Incentives

Canadian governments have offered tax-favored vehicles for saving for retirement (Registered Retirement Savings Plans - RRSPs), for the purchase of a home (Registered Home Ownership Savings Plans - RHOSPs) and for children's education (Registered

Education Savings Plans - RESPs). There is therefore, some Canadian research on the effectiveness of such vehicles.

The Canadian household saving rate diverged from its U.S. counterpart in the 1970s and remained much higher for two decades. A number of researchers attributed this divergence to the availability of RRSPs in Canada, and the absence of a comparable savings vehicle in the U.S. at that time (for example, Carroll and Summers, 1987 and Poterba et al., 1996). This seemed to suggest that RRSPs were effective in generating new savings. However the idea that RRSPs explained the divergence of the Canadian and U.S. household savings rates has subsequently been forcefully disproved by Burbidge et al. (1998) (See also Sabelhus,1997). A number of papers (for example Veall, 2001, and Milligan, 2002) have looked at the effect of taxes on RRSP contributions and participation. These studies are motivated by the idea that favoring has a larger impact on after-tax rates of return when taxes are higher. However, these studies find mixed results (for example, Milligan finds a small effect of marginal tax rates – and hence the value of tax-favoring - on participation, while Veall finds no effect of marginal tax rates on contributions.) More importantly from the point of view of this briefing, because these studies look only at RRSPs in isolation (and not other components of households' portfolios) they can't distinguish between an effect on the level of saving and an effect on allocation of saving to different assets (to RRSPs in particular). Thus a reasonable summary of the literature here might be that there is no strong empirical evidence that RRSPs have generated net new saving.

RHOSPs have been studied by Engelhardt (1996). In particular, Engelhardt studies household saving in Canadian microdata before and after the cancellation of the program. He concludes that RHOSPs did generate new saving, estimating that each dollar contributed to a RHOSP represented between 56 and 93 cents of new household saving.

RESPs are considered in a series of papers by Milligan (Milligan 2002, 2005). He documents the fact that RESPs are largely used by high-income households and concludes therefore that they do little to stimulate the savings of low-income households. HRDC (2003) also document that RESPs are largely used by high-income households. The RESP program contains a matching element, the CESG, and Milligan (2002) is particularly critical of the version the CESG that existed at that time, which he points out went largely to high-income households and so is poorly targeted and a poor use of public funds. In some ways the 2004 reform of the RESP program addressed some of Milligan's criticisms. That reform tried to increase the participation rate of low-income households by increasing matching rate of CESG and also offering a bonus (CLB) for opening an account. Milligan (2005) presents some evidence that parental expectations about their children's prospects for post-secondary education, and the fixed costs of learning about and setting up an account are factors in explaining the (continued) low use of RESPs by low-income households.

2.8 Differences between rich and poor in the response to saving incentives

A final very important point is to note that much of the economic literature (both empirical estimates and theoretical analysis) assumes that the elasticity of intertemporal substitution (the key preference parameter) is the same for rich and poor (that is, they are

equally responsive to variations in rates of return). The implication is that differences in savings behavior can arise only because of differences in budget constraints (differences in the effective rate of return faced by rich and poor households, for example). However, there are very good reasons for thinking that this is not the case, and that elasticity of intertemporal substitution differs between richer and poorer households.

The basic intuition, which is given in Browning and Crossley (2001) and further developed in Crossley and Low (2009) is that "luxuries are easier to postpone". Rich households and poor households consume different baskets of goods. The goods that make up a large fraction of the budgets of poorer households are not easily postponed. For example, eating next month is not really a good substitute for eating this month; it is fairly important to eat in both months. On the other hand, rich households consume goods that probably can be brought forward, or pushed back, in time, to take advantages of interest rates. For example, a trip to Hawaii next month (or even next year) is probably a pretty good substitute for a trip to Hawaii this month. Thus it is reasonable to think that a rich household might postpone a trip to Hawaii when interest rates are very high (so that borrowing is costly and saving advantageous) or bring a trip forward when interest rates are low (so that it is cheap to borrow, and saving offers little reward).

This intuition can be captured in a very simple model. It is very often assumed that utility is additive over time, and that utility in any period is a power function of the level of consumption in a period:

$$u(c_t) = \frac{c_t^{1-(1/\theta)}}{1-(1/\theta)}$$

This common variant of the LC/PIH has household maximize:

$$E\left[\sum_{t=1}^{T}\beta^t u(c_t)\right]=E\left[\sum_{t=1}^{T}\beta^t \frac{c_t^{1-(1/\theta)}}{1-(1/\theta)}\right]$$

With this set up one can show the following approximate relationship:

$$\Delta \ln c_{t+1}=\alpha+\theta \ln(1+r_{t+1})+u_{t+1}$$

Here $\theta>0$ is the elasticity of intertemporal substitution. It measures the effect of the rate of return (r_{t+1}) on the rate of consumption growth ($\Delta \ln c_{t+1}$). For a given θ, when r_{t+1} is larger, consumption growth is larger, meaning more saving today, and more spending in the future. For a given increase in r_{t+1} (through, for example, reducing taxes on interest income), a larger θ means a larger increase in consumption growth (and hence in saving). Note that in this formulation, the elasticity of intertemporal substitution is a constant – independent of the level of consumption or wealth. The implication is that rich and poor households can be expected to make the proportional response to a change in the (after-tax) rate of return on savings (for example, greater tax favouring). Moreover, estimates of the elasticity of intertemporal substitution from any sample (rich or poor) can be used to forecast the response of poor households to a policy designed to stimulate saving.

Now suppose we change the model just slightly, and make the (not unreasonable) assumption that utility depends on the excess of consumption over a subsistence amount $\underline{c}$:

$$u(c_t) = \frac{(c_t - \underline{c})^{1-(1/\theta)}}{1-(1/\theta)}$$

A subsistence level of consumption captures the idea that some components of consumption are very difficult to postpone or bring forward (the subsistence amount must be consumed in every period) as well as the suggestion that those components of consumption are more important for poorer households.

With this modified model one can show that the approximate relationship between rates of return an consumption growth becomes:

$$\Delta \ln c_{t+1} = \alpha + \left(\frac{c_t - \underline{c}}{c_t}\right)\theta \ln(1 + r_{t+1}) + u_{t+1}$$

And the elasticity of intertemporal substitution is now $\left(\frac{c_t - \underline{c}}{c_t}\right)\theta$, which differs between rich and poor and in particular is larger for the rich. In this setup the interest rate responses of rich or middle class families are different from the interest rate responses of poor families, and in particular, as c gets close $\underline{c}$ (that is, for households close to the minimum subsistence level they need to survive and function in society) households become completely unresponsive to rates of return. They cannot be induced to save by increasing the rate of return (through tax favouring interest etc.).

The LC/PIH is the traditional economic model of saving behavior and we have argued that it can very usefully inform thinking about the policy question at hand. At the same time, alternative frameworks for thinking about saving behavior have been garnering increasing interest in recent years. Some (but not all) of this research is

associated with "behavioral economics." In part because there is some dispute as to exactly what the term "behavioral economics" encompasses, we prefer to group the three strands of literature we discuss next under the heading "The New Economics of Saving".

3. The New Economics of Saving

3.1 Choice Architecture

There is now good evidence that some households' decisions can be affected by the framing of choices and the setting of default options. The recent popular book by Thaler and Sunstein (2008, see especially Chapter 6) provides a good summary. The advice that follows from this, that policy makers pay attention to "choice architecture", is surely sensible. Note that defaulting households into saving (with a straight forward option to opt out) is quite different from compulsion.

However, for the particular group we are considering in this review – households that are quite poor – there are unresolved issues. The first is which choices government can alter the architecture of for these households. Most of these households will not have employer sponsored pensions, for example (Lusardi, 2002), and programs based on defaulting participation, or pre-commitment to contribution increases (Thaler and Benartzi, 2004) may not be very useful for this target group.

The second issue is how the effectiveness of choice architecture varies across the income distribution. It is easy to imagine that an employed, middle class family is largely time-constrained (but not financially constrained) and so it can be defaulted into pension contributions and not opt out. It is also possible to imagine that a low income family that

is very financial constrained in the current period might be more determined to hold its saving rate down. More evidence is clearly needed.

3.2 Financial Literacy

Recent research finds that saving behavior and portfolio choices can be "explained by" - or at least are correlated with - measures of financial literacy, familiarity with financial institutions and financial planning problems, and even basic numeracy. For example, Bernheim et al., (2001b) show that children exposed to financial education have higher saving rates in adulthood; Lusardi (2001) shows that individuals learn about financial matters from parents and siblings; and Banks and Oldfield (2007) show that, in data from the English Longitudinal Study of Aging, individuals with higher numeracy are more likely to own private pensions, are more likely to invest in risky assets, have a greater understanding of retirement saving options, and report greater perceived financial security.

Lusardi (2002) suggests that financial illiteracy is acute among poor households, and perhaps particularly acute among female-headed single parent households. She also proposes that lack of experience with financial matters and institutions is a partial explanation for the low savings rates of the poor. Lusardi and others have argued that the low levels of financial literacy among poor families, in conjunction with the evidence on the importance of financial literacy in determining saving behavior, suggest an important new direction for public policy.

It is difficult to argue with this suggestion, except to note that it is only a suggestion. As Lusardi notes, low income families may be difficult to reach, and simply

mailing out booklets is unlikely to be effective. How financial literacy programs might be designed for, and delivered to, the poor, seems like a useful direction for research and policy experimentation. Proper evaluations of such programs are also needed.

3.3 Matching Programs and Changing Preferences

Finally, the last few years have seen considerable interest in matching programs as a way of stimulating saving, particularly among low income households. These programs involve special accounts, in which household contributions are matched by contributions from (usually) government. Matching rates are high – often 1:1 – implying a very high initial return to household contributions. Matched contributions are typically capped. The uses to which savings accumulated in designated accounts can be put are often proscribed (limited to educational investments or home purchases, for example).

There are several reasons that such programs are attractive to policy makers. First, as noted above, they offer much better targeting than tax-favoring. The value of the matching contributions does not rise with income (as the value of tax-favoring does) and of course participation can be limited to particular income groups. It is also the case, of course, that matching (which is often 1:1) offers very high rates of return in the short run. Beyond these considerations, though, there seems to be a hope that such programs can change behavior permanently. That is, it is hoped that such programs can raise the rate at which households save after the matching has ceased. Indeed, because these programs offer savers such high rates of return, it is hard to see how they could be cost-effective if they did not achieve some long-run change in behavior. This long run change in behavior might come about through a financial literacy effect, either because the program exposes

participants to main-stream financial institutions that they may not have otherwise encountered, or because, the program includes, in addition to matching, specific training in financial literacy. Alternatively it may be that for some individuals, experience with saving really can change preferences over the timing of consumption.

Table 1 summarizes recent evaluations of such initiatives in Canada, the U.S. and the U.K. A summary of Table 1 would be that there is some support for the proposition that these programs can raise savings rates of poorer households in the short run. However, careful studies of the effects of such programs have been undertaken only very recently, and consequently, we only have evidence on their short to medium term effects. If the real goal of such programs to change the behavior of households in a permanent way (that is to raise saving rates after the matching has ended) then we still lack important evidence on their effectiveness. Where these programs included a financial literacy training component, this does not seem to have been particularly effective, at least in the short run.

4. Conclusion

We conclude by restating the main points we have attempted to make.

1. That households with low measured current income also have low saving rates is neither surprising nor very interesting. Either measurement error in income or households' attempts to smooth consumption in the face of transitory income shocks can explain this observation. However, research has shown that Canadian households with low permanent income – households that are genuinely poor in a longer run sense – do not save very much (and save less than households who are

more affluent, in a long run sense.) This fact is worthy of attention from policy makers.

2. The fact that poor households save very little may be of concern, but it is important to note that given their lifetime income profiles and the rates of return that they face, it may in fact be optimal for them to save very little. If this is true, then *compelling* them to save more is potentially quite harmful.
3. Traditional policy leavers for increasing saving, and in particular initiatives that tax-favour saving, are unlikely to be very effective in raising the saving rates of poor households. First, tax-favouring leads a bigger change in net returns for households with higher marginal tax rates (that is, more affluent households). Second, the poor are probably less responsive to any give change in the rate of return on saving, because they spend a greater fraction of their budget on goods that hard to postpone or bring forward in time.
4. Programs that provide incentives to save through government matching households contributions' can be better targeted. Recent experiments in a number of countries provide some support for the proposition that they can raise savings rates of poorer households in the short run. However, if the real goal of such programs to change the behavior of households in a permanent way (that is to raise saving rates after the matching has ended) then we still lack important evidence on their effectiveness: longer run effects have not been assessed.
5. There is now good evidence that some households' decisions can be affected by the framing of options and the setting of default options. The advice that follows

from this, that policy makers pay attention to "choice architecture", is surely sensible. Note that defaulting households into saving (with a straight-forward option to opt out) is quite different from compulsion. However, for the particular group we are considering in this review – households that are quite poor – there are unresolved issues. The first is which choices government can alter the architecture of for these households. The second issue is how the effectiveness of choice architecture varies across the income distribution.

6. Finally, there is emerging evidence that saving behavior and portfolio choices can be "explained by" (or at least are correlated with) measures of financial literacy and even by cognitive skills. These findings open up another direction for policy development. Further research is needed in this area.

References

Alan , S. , K. Atalay, T. F. Crossley (2007), "Do the Rich Save More in Canada?" SEDAP Working Paper, no:153.

Angeletos, G., D. Laibson, A. Repetto, J. Tobacman and S. Weinberg, (2001), "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation." *Journal of Economic Perspectives*, 15(3), 47–68

Attanasio, O., J. Banks, C. Meghir and G. Weber., (1999), "Humps and Bumps in Lifetime Consumption." *Journal of Business and Economic Statistics*, 17(1), 22–35.

Attanasio, O. P. and H. Low, (2004), "Estimating Euler Equations," *Review of Economic Dynamics*, 7(2), 405-435

Attanasio, O. P and M. Wakefield, (2008), "The Effects on Consumption and Saving of Taxing Asset Returns," Prepared for the Report of a Commission on Reforming the Tax System for the 21st Century, Chaired by Sir James Mirrlees. http://www.ifs.org.uk/mirrleesreview/reports/consumption.pdf

Banks, J., R. Blundell, and S. Tanner, (1998), "Is there a retirement-savings puzzle?" *American Economic Review*, 88(4), 769–88.

Banks, J., and Z. Oldfield. (2007), "Understanding pensions: cognitive function, numerical ability and retirement saving," *Fiscal Studies*, 28(2):143-170

Bernheim, B. D., L. Forni, J. Gokhale, and L. J. Kotlikoff, (2000), "How Much Should Americans Be Saving For Retirement?" *American Economic Review,* May, 90(2), 288–92.

Bernheim, B. D., J. Skinner, and S. Weinberg, (2001), "What Accounts for the Variation in Retirement Wealth Among U.S. Households?" *American Economic Review,* 91(4), 832–57.

Bernheim, B.D., D. Garrett, and D. Maki, (2001), "Education and Saving: The long-term Effects of High School Financial Curriculum Mandates," *Journal of Public Economics,* 80:435-65.

Bernheim, B. D. and A. Rangel., (2005), "Behavioral Public Economics: Welfare and Policy Analysis with Non-Standard Decision Makers," National Bureau of Economic Research Working Paper 11518.

Browning, M. and T.F. Crossley, (2000), "Luxuries are Easier to Postpone: A Proof." *Journal of Political Economy,* 108(5), 1064-1068

Browning, M. and T.F. Crossley, (2001), "The Life-Cycle Model of Consumption and Savings," *Journal of Economic Perspectives,* 15(3), 3-22

Brzozowski, M.and Y. Lu, (2006), "Home Cooking, Food Consumption and Food Production among the Unemployed and Retired Households," SEDAP Working Paper, no:151.

Burbidge, J. B. and J. B. Davies (1994a), "Household Data on Saving Behavior in Canada," in J.M. Poterba, editor, *International Comparisons of Household Saving,* University of Chicago Press for the NBER, 11-56.

Burbidge, J. B. and J. B. Davies (1994b), "Government Incentives and Household Saving in Canada," in J.M. Poterba, editor, *Public Policies and Household Saving,* University of Chicago Press for the NBER, pp. 19-56, 1994.

Burbidge, J. B., D. Fretz, and M.R. Veall, (1998), "Canadian and American Saving Rates and the Role of RRSPs," *Canadian Public Policy*, 24(2), 259 - 263.

Carroll, C., L. Summers, (1987) "Why Have Private Saving Rates in the US and Canada Diverged?" *Journal of Monetary Economics,* 20(2),249–279

Carroll, C., (1997), "Buffer-stock saving and the life-cycle/permanent income hypothesis", *Quarterly Journal of Economics,* 112(1), 1–55

Carroll, C., (2001) "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)", *Advances in Macroeconomics*, 1(1), 1003-1003.

Crossley, T.F. and H. Low, (2009) "Is the Elasticity of Inter-temporal Substitution Constant?", forthcoming, *Journal of the European Economics Association.*

Deaton, A., (1991), "Saving and liquidity constraints," *Econometrica,*, 59(5), 1221–48.

Duflo, E., W. G. Gale, J. Liebman, P. Orszag and E. Saez, (2006), "Savings Incentives for Low – and Middle-Income Families: Evidence from a Field Experiment," *Quarterly Journal of Economics,* November 2006 1311-1346.

Duflo, E., W. G. Gale, J. Liebman, P. Orszag and E. Saez, (2007), "Savings Incentives for Low- and Moderate-Income Families in the United States: Why is the Saver's

Credit Not More Effective?" *Journal of the European Economic Association* 5(2-3), 647-661

Emmerson, C., P. Harvey, R. Madden, N. Pettigrew, G. Tetlow, M. Wakefield, (2007) "Final Evaluation of the Saving Gateway 2 Pilot: Main Report," published on HM http://www.hm-treasury.gov.uk/d/savings_gateway_evaluation_report.pdf

Engelhardt, G. V., (1996), "Tax Subsidies and Household Saving: Evidence from Canada," *The Quarterly Journal of Economics,* 111(4), 1237-1268.

Engelhardt G. V. and A. Kumar, (2007), "Employer matching and 401(k) saving: Evidence from the Health and Retirement Study," *Journal of Public Economics* ,91, 1920–1943

Engen, E. M, W. G. Gale and J. K. Scholz, (1996), "The Illusory Effects of Saving Incentives on Saving," in "Symposia: Government Incentives for Saving," *The Journal of Economic Perspectives,* 10(4),113-138.

Engen, E. M., W. G. Gale, C. E. Uccello. (1999), "The Adequacy of Household Saving," *Brookings Papers on Economic Activity,* 2, 65-165.

Friedman, M., (1957), *A Theory of Consumption Function,* Princeton University Press.

Gourinchas, P.-O., and J. A. Parker, (2002), "Consumption over the life-cycle," *Econometrica,* 70(1), 47-89.

HRDC, (2003), "Formative Evaluation of the Canada Education Savings Grant Program," Evaluation and Data Development, Strategic Policy, Human Resources Development Canada, Ottawa.

Hubbard, R. G., J. Skinner and S. Zeldes. (1994), "Precautionary Savings and Social Insurance,." *Journal of Political Economy,* 103, 360– 99.

Hugggett, M. and G. Ventura, (2000), "Understanding why high income households save more than low income households," *Journal of Monetary Economics* 45, 361-397

Hurd, M. D, and S. Rohwedder, (2006), "Some Answers to the Retirement–Consumption Puzzle," National Bureau of Economic Research Working Paper 12057(February)

Kempson, E., S. McKay and S. Collard (2003), "Evaluation of the CFLI and Saving Gateway Pilot Projects," published on HM Treasury website: http://www.hm-treasury.gov.uk/d/SG_report_Oct03_162.pdf

Kempson, E., S. McKay and S. Collard (2005), "Incentives to save: Encouraging saving among lower-income households. Final report on the Savings Gateway pilot project," published on HM Treasury website: http://www.hm-treasury.gov.uk/d/incentivestosave_150305.pdf

Kempson, E., A. Atkinson and S. Collard (2006), "Saving for children: A baseline survey at the inception of the Child Trust Fund," http://www.pfrc.bris.ac.uk/publications/pensions_savings/Reports/Saving_for_children_report.pdf

Kingwell, P., M. Dowie, B. Holler, C. Vincent, D. Gyarmati and H. Cao , (2005). "Design and Implementation of a Program to Help the Poor Save: The learn$ave," published on SRDC site: http://www.srdc.org/uploads/learnsave_implementation.pdf

Kusko, A, J.M. Poterba and D. Wilcox ,(1998), "Employee Decisions with Respect to 401(k) Plans: Evidence from Individual-Level Data," in Olivia Mitchell and Sylvester Schaefer, eds., Living with Defined Contribution Pensions: Remaking Responsibility for Retirement. (University of Pennsylvania Press: Philadelphia).

LaRochelle-Côté, S., J. Myles and G. Picot, (2008), "Income Security and Stability During Retirement in Canada," Analytical Studies Branch Research Paper Series, Statistics Canada.

Leckie N., M. Dowie and C. Gyorfi-Dyke, (2008). "Learning to Save, Saving to Learn: Early Impacts of the learn$ave Individual Development Accounts Project," published on SRDC site: http://www.srdc.org/uploads/learnsave_full_EN.pdf

Lusardi, A., (2000), "Explaining Why So Many Households Do Not Save," Mimeo.

Lusardi, A., (2002), "Increasing Saving among the Poor: The role of Financial Literacy," *Poverty Research News*.

McBride, M. A., M. Lombe and S. G. Beverly, (2003), "The effects of individual development dccount programs: perceptions of participants." *Social Development Issues*, 25(1/2), 58-73

Mills G., R. Patterson, L. Orr and D. DeMarco, (2004), "Evaluation of the American Dream Demonstration," Abt Associates Inc. Cambridge, MA.

Mills, G., W. G. Gale and R. Patterson, (2005), "Effects of Individual Development Accounts on Household Saving Behavior: Evidence from a Controlled Experiment," Abt Associates and the Brookings Institution.

Milligan, Kevin, (2002), "Tax Preferences for Education Savings: Are RESPs Effective?" Commentary 174, C.D. Howe Institute. http://www.cdhowe.org/pdf/commentary_174.pdf

Milligan, K., (2002), "Tax-preferred Savings Accounts and Marginal Tax Rates: Evidence on RRSP Participation," *Canadian Journal of Economics,* 35, 436-456

Milligan, K., (2005), "Lifecycle Asset Accumulation and Allocation in Canada," *Canadian Journal of Economics,* 38(3), 1057-1106.

Milligan, K., (2005), "Who uses RESPs and why," in Charles M. Beach, Robin W. Boadway and R. Marvin McInnis (eds.) *Higher Education in Canada*, McGill-Queen's University Press, 467-494.

Modigliani, F., 1976, "Life-cycle, individual thrift, and the wealth of nations," *American Economic Review*, 76(3), 297–313

Modigliani, F. and R.H. Brumberg, (1954) "Utility analysis and the consumption function: an interpretation of cross-section data," in Kenneth K. Kurihara, ed., *Post-Keynesian Economics*, New Brunswick, NJ. Rutgers University Press. 388–436.

Papke, L. and Poterba J.M. (1995), "Survey of Evidence on Employer Match Rates and Employee Saving Behaviour in 401(k) Plans," *Economic Letters* , 49, 313-317

Poterba J., S.Venti, and D. Wise, (1996), "How Retirement Saving Programs Increase Saving", (in "Symposia: Government Incentives for Saving"), *The Journal of Economic Perspectives,* 10(4), 91-112

Sabelhaus, J., (1997), "Public Policy and Saving in the United States and Canada," *Canadian Journal of Economics,* 30, 253-275

Schillington, R., (2003), "New Poverty Traps: Means-testing and Modest-Income Seniors," *C.D. Howe Institute Backgrounder No. 65, C.D. Howe Institute.* http://www.cdhowe.org/pdf/backgrounder_65.pdf

Scholz, J. K., A. Seshadri, and S. Khitatrakun, (2006), "Are Americans Saving 'Optimally' for Retirement?" *Journal of Political Economy*, 114(4): 607–43

Skinner, J., (2007), "Are You Sure You're Saving Enough for Retirement?" *Journal of Economic Perspectives,* 21(3), 59–80

Statistic Canada Report, (2006) "The wealth of Canadians : An Overview of the Results of the Survey of Financial Security 2005"

Thaler, R., S. Benartzi, (2004), "Save more tomorrow™: using behavioral economics to increase employee saving," *Journal of Political Economy,* 112(1), S164–87.

Thaler, R.H., and C.R. Sunstein, (2008), *Nudge: Improving Decisions About Health, Wealth and Happiness.* Yale University Press.

Veall, M.R. (2001) "Did tax flattening affect RRSP contributions?" *Canadian Journal of Economics* 34, 120-31

Table

Table 1: Matching Saving Programs

Country	Program	Description	Results
Canada	**Learn$ave**	A randomized control trial designed to test the effects of matched savings incentives on low income adults. Each dollar of savings by participants in the treatment groups is matched by a three dollar credit over a 3 year period (to a maximum of $250 per month and $1,500 over the three-year period). Some treatment groups also received training in financial management.	Preliminary results are reported in Leckie et. al, (2008). By month 18, the treated group receiving matching funds only had saved 71 per cent more than the control group. Treatment groups are more likely to maintain a household budget, but financial management training has had little effect on savings.
Canada	**Canada Education Savings and Grants (CESG) & Canada Learning Bonds(CLB)**	The *CESG* gives a matching grant of 40 %, 30% or 20% of the value of registered educational savings plan (RESP) contributions to low, moderate, high income households respectively. The *CLB* is a further effort to encourage low income households to set up RESP accounts. It offers a $500 dollar grant to any child born after 2003 in a low income household (plus $100 each year to age 15).	Preliminary analysis of CLB reported in Knight et al, (2008) shows that out of 481,953 eligible children only 38,268 have received CLB since the program initiated, implying a takeup rate of only 7.9 percent. Although this number has risen from 4.6% (in 2006), it is still very low.
U.S.	**H&R Block Retirement Saving Experiment**	This is a large scaled randomized field experiment designed to estimate the effect of matching rates on the willingness of low and middle income families to contribute to IRAs.	Duflo et.al. (2006) show that: i-) Matching rates have an effect on take up rates ii-) Matching rates also have a strong effect on overall contributions iii-) Effect of matching is significant for all income groups but increases

Continues on the next page

			with income. iv-) Professional tax assistance, and information provision, play an important role in encouraging savings
U.S.	**Saver's Credit**	A U.S federal government program, this credit provides a federal income tax reduction of up to 50% of funds contributed to a 401(k) or IRA.	Duflo et. al.(2007) find a significant, but very modest, effect on contributions to IRAs. The authors also compare this effect and that of the H&R match experiment. They note that contributions to Saving Credit are smaller than the H&R experiment, even though the effective match provided by the Saver's credit is larger for lower income households.
U.S.	**401(k)**	Many employer sponsored 401(k) programs involve matching of employee contributions.	There is mixed evidence on the effectiveness of matching. Papke and Poterba (1995) finds i-) a positive relationship between match rates and participation rates and ii-) a weak link between match rates and the value of employee contributions. Kusko , Poterba and Wilcox (1998) find that the existence of matching raises 401(k) savings, but that higher matching rates do not lead to further increases in savings. Engelhardt and Kumar (2007) find that the 401(k) saving response to matching is quite inelastic (and so conclude that matching is a rather poor policy instrument to raise retirement savings).

Continues on the next page

U.S.	**IDAs (The Tulsa Experimental IDA)**	Individual development accounts (IDAs) are saving accounts that provide low income households with matching of contributions. Withdrawals from these accounts must be used for special purposes, such as a home purchase, business start-up, or investment in education.	Mills et.al (2004) evaluates this first controlled field experiment (Tulsa) on IDAs. They report i-) an increase in the savings of the low income groups ii-) savings accumulated in IDAs is used for the intended purposes. Mills et.al (2006) show that the IDA program had different effects across sub-groups. Particularly for African Americans (who are more likely to face discrimination in the lending market) the IDA program raised the homeownership. Therefore the IDA program was also effective in reducing financial barriers to home ownership among African-Americans. McBride et al. (2003) argue that IDAs have positive physiological effects.
UK	**Savings Gateway (SG 1-2)**	A saving account that is designed to stimulate low income households savings by offering matching contributions. The first pilot (SG1) offered 1:1 pound matching up to a £25 monthly contribution limit and an overall cap of £375. The second pilot (SG2) was initiated in 2005. It offered alternative match rates (£1 Government contribution per every £1, £2 or £5 saved) and different monthly contribution limits (£25, £50, £125) to participants. The second pilot was also tested on a far wider income group.	*Kempson et al. (2003)* report that i-) the program was successful in encouraging low income adults to save, ii-) Most of the savings are new savings, iii-) most participants declared that they intended to continue to save regularly after their accounts matured. *Kempson et al. (2005)* follows up this short run evaluation of SG1 with an evaluation of medium term outcomes. After 3 (or more) months after the accounts matured, i-) Most participants (91%) still had their savings accounts, ii-) Nearly half of them (41%) still depositing regularly, and iii-) 32% of participants said that they were more likely to plan for retirement. *Emmerson et al. (2007)* evaluate the SG2. They

Continues on the next page

			report that i-) the effects of pilot varied across income groups, ii-) for lower income participants, SG2 may have lead to new savings, iii-) For those with higher incomes, SG2 savings are associated with the reduction in their balances on other assets, and iv-) There is also evidence that the program also brought some individuals into contact with mainstream financial institutions for the first time.
UK	**Child Trust Funds (CTF)**	Not quite a matching program, the CTF provides £250 for every child born after August 2002. An additional £250 deposited for low income families. Children will also receive similar payment at age 7.	Kempson et al. (2005) report a preliminary analysis of this program. They find that participation is high (70 % percent of children had either a savings account or some form of investment.)

Figure

Figure 1: Saving Rates across Income Groups

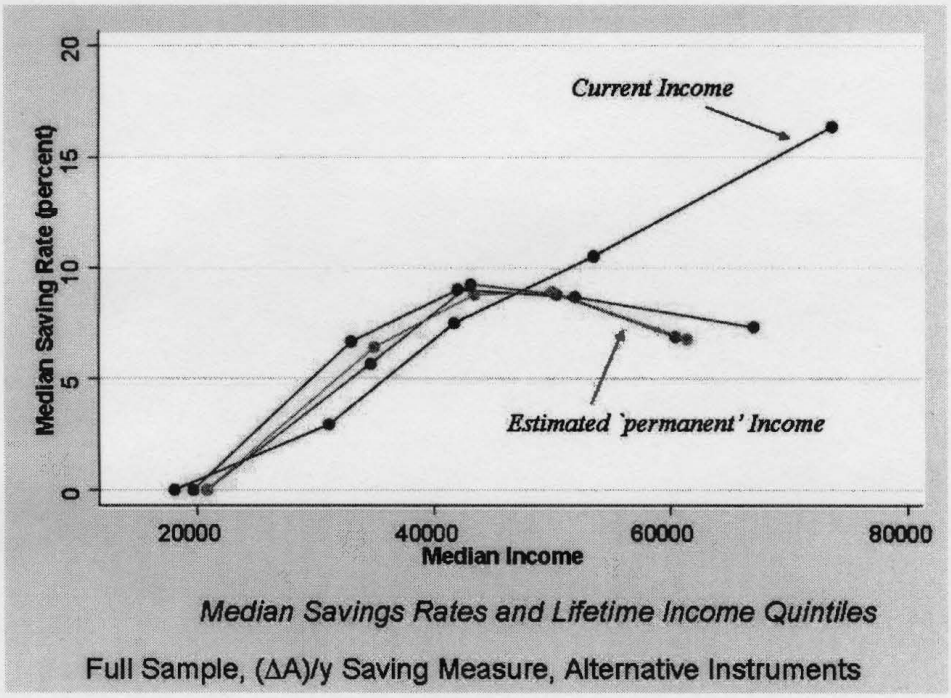

Note to Figure 1:

- Adapted from Alan, Atalay and Crossley (2007)

VII

Conclusion

This thesis is composed of an introduction and five independent chapters. The first chapter after the introduction investigates the living standards of Canadian retirees and the adequacy of their financial preparations for retirement. The third chapter revisits a long standing question of whether households with higher *lifetime* income save a larger fraction of their income. The fourth chapter investigates the effect of taxation on households' portfolio allocation decisions. The fifth chapter investigates the problems associated with the estimation of intertemporal allocation parameters via linearized Euler equations. The chapter before the conclusion reviews the economic literature on the saving behavior of low income households.

The first chapter after the introduction contributes to the literature on the retirement saving adequacy. We use an alternative source of information that has not been employed in Canada in this context before. We explore the responses of Canadian retirees to subjective survey questions administered in General Social Surveys and in the 1975 Retirement Survey. Our results differ markedly from standard analyses of consumption and income. We find a significant portion of Canadians report being at least as satisfied with their finances as they were in the year prior to retirement. Our results also suggest a strong positive relationship between involuntary retirement and financial dissatisfaction. These findings underline the necessity of new panel datasets that contain

income and expenditure information a long with a range of other living standard measures (including subjective ones) around retirement. Further analysis using such datasets would help us to better understand retirement saving behavior.

The third chapter contributes to the literature on the relationship between lifetime incomes and saving rates. Our empirical analysis shows that except for poorest households, household saving rates do not differ substantially across life time income groups. This finding implies that standard economic models of saving (which by and large imply constant saving rates by lifetime income) might provide reasonable guidance to certain policy questions. On the other hand our finding that poor households' saving rates are different than the rest of the population provides a useful guide for future research priorities.

The fourth chapter contributes to the literature on the effect of taxation on portfolio choice. We use an alternative and previously unexploited source of variation in tax rates for identifying tax effects on portfolio allocation. Our results show a statistically significant but economically modest effect of taxation on households' portfolio choices. Hence, we find an increase in tax rates may cause households to hold more tax favored saving assets. The next step in this line of research could focus on understanding the dynamics of portfolio adjustments. Most studies of household portfolios investigate the effect of current tax rules on current household asset holdings. If portfolio rebalancing involves significant transaction costs, adjustments to portfolios will not be continuous. One fruitful avenue of inquiry could be to study the cumulative effect of previous years' tax systems on asset allocations of households.

The fifth chapter contributes to the literature on the estimation of intertemporal allocation parameters via linearized Euler equations. Our findings show that problems associated with linearized Euler equations are strongly related to assumed economic environments. In addition, we also find a tradeoff between validity and relevance of conventional instruments used in the estimations. These findings imply that there are environments (parameter values) where linearized Euler equations work well and environments in which instrument problems plague estimation. These results reconcile the differing conclusions present in this literature. Our finding, that one is more likely to extract the response to interest rates from the pool of wealthy households and extract the response to uncertainty from the pool of buffer stock savers, suggests potential avenues for future research. A starting point may be to study, through simulation, how a subsistence level of consumption (a consumption floor framework) affects the estimation of intertemporal allocation parameters for different groups.

The last chapter before the conclusion contributes to the literature on the saving behavior of low income households. This chapter reviews the current state of research on saving and focuses on the implications of these studies for low income households. We outline the motives shaping the saving behavior of low income households and review evidence on the efficacy of a variety government programs designed to stimulate saving among this group. Further progress on this issue requires more empirical evidence and new datasets that contain comprehensive information on low income households. In a Canadian context, there are government programs (for example, the Canada Learning Bond) which could be evaluated empirically with the help of new Canadian datasets.