RRSPs and National Saving in Canada

Christopher Ragan
Department of Economics
McGill University

December 1995

This paper was prepared for the Economic Studies and Policy Analysis Division of the federal Department of Finance under contract No. A506046. The views expressed here are those of the author and do not necessarily reflect the views of the Department of Finance.

RRSPs and National Saving in Canada

Table of Contents

RRSPs and National Saving in Canada

I. Saving, Growth and Policy

This paper examines the flow of national saving and how that flow is influenced by tax policy—specifically, tax policy aimed at encouraging the flow of private saving. In Canada, there are two major tax instruments of this type: Registered Retirement Saving Plans (RRSPs) and employer-sponsored Registered Pension Plans (RPPs). Before delving into the details, however, we first explore the motivation behind this paper. Why do we care about national saving? And why might there be a belief that government policy in this regard is desirable or can be effective?

A. Importance of National Saving to the Economy.

Investment in physical capital has long been viewed as an important determinant of economic growth. Research which emphasizes the role of capital accumulation dates back at least to Solow (1956); a recent review—suggesting that the statistical link between investment and growth is one of the few robust relationships related to growth—is provided by Levine and Renelt (1992). Investment has been a traditional focus partly because increases in the stock of physical capital directly result in higher future production and income, but also because the process of investment is one method of bringing technological advances into the workplace—better technologies are often embodied in new plant and equipment.

If the role of investment in promoting economic growth is relatively uncontroversial, the connection between a country's flows of saving and investment has certainly been subject to debate. At the root of this debate are competing beliefs about the extent to which the capital market of any particular economy is integrated into the world capital market. Thus, one way to organize our thoughts about the saving-investment relationship is to think separately about closed and open economies.

• Closed Economies. In a closed economy, the only way that investment can occur is through saving; private agents give up current consumption in order to increase future

consumption. Thus saving and investment are exactly equal in a closed economy, even though saving and investment are fundamentally different actions, undertaken by different economic agents. Saving is typically done by households who choose to defer current consumption; investment is typically undertaken by firms deciding to expand productive capacity.

The equality of saving and investment in a closed economy is obviously not a coincidence—it results from an operative price mechanism. Any excess supply of loanable funds leads to a fall in real interest rates until the market for loanable funds clears. When this market clears, the quantity of funds demanded for investment purposes equals the quantity of funds supplied by savers.

One implication of the saving-investment equality in a closed economy is that if an objective of government policy is to increase the amount of domestic investment (in an effort to promote economic growth), then policy can sensibly be aimed at *either* side of the market. Policy can be designed to directly stimulate domestic saving or to directly stimulate domestic investment. A policy that directly encourages domestic saving will lower real interest rates and thereby increase the amount of domestic investment. Alternatively, a policy that directly encourages domestic investment will raise real interest rates and thereby induce domestic residents to increase their saving. In both cases, the equality of domestic saving and domestic investment is an equilibrium outcome.

• Conventional View of Open Economies. Things are quite different in open economies when capital is fully mobile across international boundaries. In such a setting, there is no reason in principle why saving and investment need to be equal or even closely related. In an open economy, an excess of domestic saving over domestic investment can be accommodated by lending the excess to foreigners—a current account surplus. Conversely, an excess of domestic investment over domestic saving can be financed by borrowing from foreigners—a current account deficit.

What might therefore be called the "conventional view" is that in open economies with high degrees of capital mobility, there need be no close relationship between domestic saving and domestic investment. Moreover, if a policy objective is to increase the amount of domestic investment, then it matters which side of the market the policy aims at—directly encouraging domestic saving will have quite different effects from directly encouraging domestic investment.

In a small country which faces a given world real interest rate, a policy which directly encourages domestic investment will have little or no effect on the amount of domestic saving, but will lead to an increase in the current account deficit. In contrast, a policy which directly encourages domestic saving will have no effect on the amount of domestic investment, but will lead to an increase in the current account surplus.

Though an increase in saving in an open economy may not have any effect on the rate of domestic capital formation, it will affect who owns that capital. This is because the separation of saving and investment in open economies also implies a separation of the *location* of investment from the *ownership* of the investment. In a closed economy, the investment which gets undertaken domestically is necessarily financed by the saving of domestic residents, and thus it is owned by domestic residents. But this is not the case in an open economy. If an increase in domestic investment occurs without an equivalent increase in domestic saving—and thus there is an increase in the current account deficit—then the increase in domestic investment is only possible by increasing net foreign indebtedness. Thus, while there will indeed be more productive capacity in the domestic country, not all of the accompanying higher production can be consumed domestically; some will have to be remitted to the foreigners (either in interest or in profits) who financed the investment in the first place.

Thus, while an increase in saving in an open economy may not increase the amount of domestic capital formation, it will increase the share of future production which can be consumed by domestic residents. In other words, while an increase in domestic saving may have no effect on domestic future GDP, it will increase domestic future GNP. To the extent that we are interested not in how much is *produced* domestically, but rather in how much can be *consumed* domestically on a long-term basis, an increase in domestic saving may therefore be a good thing.[1]

• An Emerging View of Open Economies. This conventional view of the clear separation of saving and investment in open economies has come under attack in the last decade, beginning with the landmark study by Feldstein and Horioka (1980). They showed, using a sample of several countries over time, that there is a strong positive correlation between a

[1] There is no free lunch here, however. An increase in national saving does imply higher future consumption possibilities, but this is only possible by reducing current consumption.

country's saving rate and investment rate. They interpreted this correlation as evidence that capital is not highly mobile across international boundaries. Though others have questioned the Feldstein-Horioka interpretation of this high correlation, the basic fact itself appears to be quite robust to different sample periods and different countries.

The Feldstein-Horioka finding suggests that even though there may appear to be few restrictions placed on the international mobility of capital, it may still be the case that the owners of capital have a preference for conducting transactions in relatively local capital markets. (One possible reason for this type of locational preference is that firms finance a disproportionate amount of their investment out of retained earnings, rather than by expanding their bank debt or by issuing debt or equity.) Thus, even though economies are open and much capital moves across international boundaries, many countries behave almost *as if* they are closed in the sense that there is a strong correlation between domestic saving and domestic investment.

If capital really is as immobile as the Feldstein-Horioka results suggest—for whatever reason—then public policy in an open economy can look very much like public policy in a closed economy. That is, if government wants to design policies which will increase the amount of domestic investment, then policies which directly encourage domestic saving may indeed have some effect.

B. Recent Concern over Deficient Saving.

The foregoing discussion suggests that policies designed to directly encourage the flow of national saving may be successful in promoting domestic investment and, through that channel, in increasing future economic growth. Such an increase in national saving generates obvious benefits for future generations, but this occurs at the cost of reducing the current generation's consumption. A policy thus designed to increase the current flow of saving (and thus to reduce current consumption) could be justified if there were some basis for believing that the economy in the absence of such policy would be saving too little relative to what is socially optimal.

One reason that the private economy may save less than what is optimal is that the current tax system distorts private decisions. For example, the taxation of capital income lowers the after-tax real interest rate and thus provides a disincentive for individuals to save. A policy

to reverse the effects of this distortion may thus help return the economy to its appropriate level of saving. This distortion, and policies to eliminate it, is discussed in detail in Section 2.

Aside from the belief that existing distortions may result in too little saving, there are two other often-heard reasons for being concerned about low national saving rates. The first is a short-term concern and is related to the ability of countries to run persistent current account imbalances. For example, Feldstein (1995) has recently argued that an important underlying cause of the 1994 collapse in the Mexican Peso was Mexico's reliance on foreign capital to finance its investment—that is, the problem was fundamentally one of deficient Mexican saving. Feldstein's argument is surely controversial; less controversial is that the degree of capital (im)mobility and the level of Mexican saving were not the *only* cause of the Mexican problem; the exchange-rate regime and the degree of monetary expansion certainly contributed to the December 1994 Peso crisis.

The second concern over low national saving rates is a long-term concern and is related to the issue of economic growth. As discussed above, a higher national saving rate implies (with limited capital mobility) a higher domestic investment rate which in turn leads to higher future economic growth. This long-term concern is the one more relevant to Canada and other developed countries.[2]

The concern over low national saving rates probably has its origins in the now well-documented "productivity slowdown" which began in the early 1970s. Recent anxiety—more in the United States than in Canada[3]—has surely been heightened by the emergence of countries like the Asian Tigers (Hong Kong, Singapore, South Korea and Taiwan) which continue to display both high growth rates and high saving rates. But it is relatively easy to console ourselves when confronted with the experience of the Asian Dragons; after all, they are developing countries, and periods of high growth and constrained consumption in such countries

[2] Keep in mind, however, that people who argue in favour of promoting national saving on the grounds that it raises future consumption possibilities are essentially arguing in favour of an intertemporal redistribution away from the current generation and toward future generations.

[3] See Bosworth (1989) and Nordhaus (1989) for concerns over the low U.S. saving rate.

is expected. It would be virtually impossible for those countries to sustain indefinitely their high growth rates of the recent past.[4]

We might dismiss the evidence on the saving-growth relationship in developing countries as being largely irrelevant to a country like Canada, but it is harder to dismiss similar evidence from developed countries. Figure 1, for example, plots the growth rate of real per capita GDP against net national saving as a share of GDP. Each point corresponds to a country and decade; the sample includes the G7 countries over the 1960s, 1970s and 1980s. Each point therefore shows the average annual per capita growth rate in that decade and the average annual net national saving rate. The line shows the predicted values from a simple linear regression of the growth rate against (a constant and) the saving rate.[5]

Figure 1 obviously shows a very strong correlation between saving rates and per capita growth rates. Though this evidence is clearly suggestive of a link between saving and economic growth, it is important to keep in mind that such correlations do not provide convincing evidence of a causal relationship. Indeed, Carroll and Weil (1994) have recently argued that this positive correlation between saving and growth may actually have its source in a reverse causality—that is, it is the high growth rates which cause the high saving rates rather than vice versa. This argument should not be surprising. Even very simple models with overlapping generations of life-cycle agents show this result: National saving would be increasing in the population growth rate and in the rate at which each new generation is wealthier than the previous one.

Such evidence as that shown in Figure 1 naturally leads to calls—justified or not—for the design and implementation of government policies to encourage both domestic investment and domestic saving. Individual Retirement Accounts and 401(k) plans in the United States, and RRSPs and RPPs in Canada, are examples of large and important policy initiatives which have

---

[4] Paul Krugman (1994) makes precisely this argument in a well-known article in *Foreign Affairs*. He also notes that not all of the Asian Dragons have similar types of growth. Singapore's growth, for example, has relied very heavily on the growth of inputs, whereas Hong Kong shows relatively much more growth in total factor productivity. Some of Krugman's arguments are based on earlier work by Alwyn Young (1992).

[5] The regression and the figure excludes Japan during the 1960s. For that decade, the Japanese saving rate was 25.6% and the average annual growth rate was 9.5%, putting that country-decade point far in the North-east corner of Figure 1 (and well above the regression line). This data is drawn from De Gregario (1993), Tables 2 and 4.

been designed to offer tax assistance to saving, thereby providing an incentive for individuals to increase saving for their retirement.

C. Basic Description of RRSPs and RPPs.

The Canadian government introduced Registered Retirement Saving Plans (RRSPs) in 1957, permitting individuals to make deposits into saving plans which receive special tax advantages. Similar tax advantages were already available to some individuals through employer-sponsored Registered Pension Plans (RPPs) and Deferred Profit Sharing Plans (DPSPs). The source of the tax advantage is that contributions to RRSPs, RPPs and DPSPs are tax deductible in the year they are made; furthermore, the funds accumulate within the plan without taxation of interest or dividend income. The benefit to individual contributors from such tax-deferred saving plans is reflected by the very significant growth in annual contributions since their inception.

Figure 2 shows the annual flow of contributions into RRSPs and RPPs since 1960. Though it is not clearly seen in the figure, RPPs were far more important than RRSPs in the early 1960s. Given that RRSPs had only been created in 1957, this should not be surprising. For example, individuals contributed $27.7 million into RRSPs in 1960s, but contributed over $330 million to RPPs in that same year. By 1970, RRSPs had become relatively more important, but were still dominated by RPPs; in that year, individuals contributed $225 million to RRSPs and almost $730 million to RPPs. But in 1976, annual contributions to RRSPs overtook those to RPPs, and the gap has increased dramatically since then. By 1992, the annual flow of contributions to RRSPs totalled almost $15 billion, whereas only $6.7 billion was being contributed to RPPs.[6]

It is clear that RRSPs have evolved to the point today where the annual contributions represent an enormous sum of money. Note also that the dramatic increase in annual contributions is not just reflecting a steadily rising price level over this period; the RRSP contributions in Figure 2 are expressed in constant dollars. Over the thirty years shown, real annual contributions to RRSPs fell only once, at the beginning of Canada's most recent recession in 1990.

[6] This data is drawn from *Taxation Statistics*, various years. Note that employers' contributions to RPPs are not included in Figure 2.

The growth of RRSPs as a vehicle of individual saving probably reflects two factors. One is the increased public awareness of the existence of RRSPs. The other is the increased generosity of RRSPs. (These are unlikely to be independent of each other.) Table 1 shows how the contribution limits to RRSPs have evolved since their inception in 1957. The generosity of RRSPs increased dramatically in 1972, when the maximum contribution increased from 10% to 20% of earned income. Since that time, the adjustments to the RRSP program have mostly taken the form of increasing the contribution limit to reflect increases in the overall price level. For example, between 1972 and 1994, the maximum contribution as a share of earned income fell from 20% to 18%, but the absolute upper limit increased from $4,000 to $13,500. Another important change in the RRSP program, which took place in 1991, was an attempt to better harmonize the limits for individuals with employer-sponsored pension plans to the limits for individuals with no employer-sponsored plans. This was accomplished by introducing the concept of the Pension Adjustment (PA).[7]

D. Outline of this Paper.

The purpose of this paper is to examine the relationship between RRSPs and the level of national saving. The basic question we examine is: Do RRSPs increase the level of national saving?

The paper is organized as follows. The next section examines some basic theoretical issues related to consumption, saving and RRSPs. The purpose is simply to lay out a framework for thought—one that both is an aid for thinking about the various issues and also helps in the interpretation of the empirical results which follow. Section 3 offers a very brief examination of aggregate saving data in Canada. Over the past twenty-five years, there has been a clear (though not smooth) downward trend in national saving. The combined saving of the private and public sectors is currently about 4 percent of GNP. Section 4 then discusses in some detail the empirical study by Sargent (1995) which uses a large cross-section of households to model saving behaviour. Sargent's approach follows closely that used in the United States by Gale and Scholz (1994), and finds broadly similar results. Specifically, Sargent finds that reductions in

[7] See Horner and Poddar (1992) for a clear discussion of pension reform in Canada, and a detailed discussion of the Pension Adjustment.

the generosity of the RRSP program can be expected to generate substantial increases in government and national saving. Section 5 discusses Sargent's results, and attempts to provide an understanding based on the theoretical discussion in Section 2. The issue of Ricardian Equivalence is also addressed. Section 6 offers some final remarks, and points to other issues that deserve examination, such as the impact of RRSPs on the distribution of income.

II. Central Analytical Issues

This section is divided into five parts. The first briefly examines the main elements of a simple life-cycle model of consumption—the framework which provides the starting point for most research on consumption and saving. The second part adds RRSPs to this simple framework. The third and fourth parts examine, respectively, the effect of RRSPs on private saving and public saving. The final part then combines the results from the preceding two parts to examine the relationship between RRSPs and national (public plus private) saving.

A. The Simple Life-Cycle Model.

We use the simple life-cycle model of consumption as the basic framework of analysis for examining the relationship between RRSPs and an individual's private saving. It thus plays a central role in the overall examination of RRSPs and national saving. Despite the fact that this model has recently come under attack for not capturing some of the real-world determinants of individual saving behaviour (e.g., Hubbard et al 1994; Thaler 1994), its value is immense—and not just for its simplicity. Perhaps more important for our purposes is this model's pervasiveness in the minds of both policy makers and economists. Understanding the central insights of the life-cycle model is thus central to understanding the design of many current policies.

The starting point for the life-cycle model is to posit an individual who lives over several periods, with some known pattern of income over those periods. Typically, a given individual will have low income early in life, higher income as he reaches the peak of his career, and then lower income during retirement. Furthermore, the individual has access to a capital market in which he can borrow or lend at the real interest rate. This is all taken as given in the simple life-cycle model.

The question in this setting is then: What determines this individual's pattern of consumption (and therefore saving) over the life cycle? The central insight of the life-cycle approach to consumption is the importance of what is called "consumption smoothing"—which reflects both the *desire* and the *ability* of the individual to consume at more-or-less constant levels across the life cycle, even though the life-cycle pattern of income may be very uneven.

The *desire* for consumption smoothing comes from the typically-assumed convexity of the individual's preferences between current and future consumption. Convexity of preferences simply reflects the assumption that the individual would prefer to consume "medium" amounts both currently and in the future rather than "very high" amounts now (later) and "very low" amounts later (now). The *ability* for the individual to achieve such consumption smoothing comes from his assumed access to a well-functioning capital market in which he can borrow or lend at the going real interest rate. It is this access to the capital market which permits the individual to lend when income is greater than desired consumption and to borrow when income is less than desired consumption.

• Predictions of the Model. Consumption smoothing—the central feature of the life-cycle model—gives rise to a very strong prediction regarding which factors will affect the individual's consumption and saving behaviour. Since the individual uses the capital market to spread his lifetime stream of income across the years to conform with his desired pattern of consumption, there are two variables which are sufficient to determine the individual's time-path of consumption. The first is the individual's wealth, the discounted present value of the lifetime stream of income. The second is the real interest rate, the relative price of current consumption in terms of forgone future consumption.

Given these two variables, the time-path of consumption is uniquely determined. One important point made by the life-cycle model is that changes in current income *that do not affect the individual's wealth* have no effect on the individual's time-path of consumption (though they obviously then do affect his saving since the flow of saving is simply equal to income minus consumption).

Obviously, however, changes in wealth or in the real interest rate typically lead to changes in the individual's consumption pattern. For example, an increase in income in any year that is not offset by an equal reduction in some later year translates into an increase in the

individual's lifetime resources—that is, an increase in wealth. If both current and future consumption are normal goods, then such an increase in wealth will increase consumption now and in the future. The current flow of saving will also increase in this case (some of the increase in current income is saved so that future consumption can be increased).

An increase in the real interest rate also changes the time-path of consumption. It not only reduces measured wealth (because future income gets discounted at a higher rate), but it also changes the intertemporal relative price. Current consumption is now more expensive in terms of foregone future consumption. This introduces both income and substitution effects. The substitution effect unambiguously leads to lower current consumption; the direction of the income effect depends on whether the individual is a net debtor or a net creditor. Thus, we are generally unable to predict what happens to an individual's consumption when there is an increase in the real interest rate. This reflects more than just our ignorance about the direction of the income effect; even if we knew the direction, the income and substitution effects may well be pushing in opposite directions, in which case the overall effect depends on their respective magnitudes.

• Adding Taxes to the Model. By adding a proportional income tax to the basic life-cycle model, a distortion is generated. If the income tax applies to labour and interest income equally (and by symmetry exempts interest payments from taxable income), then the after-tax interest rate is clearly lowered. This lowers the intertemporal relative price, and generates a clear substitution effect toward more current consumption and away from current saving.[8]

Adding such an income tax to a model which has no taxes is not a particularly interesting policy experiment. Even if it were, there would be no clear prediction about the effect of the tax on current consumption—the substitution effect would increase consumption but the income effect would reduce it. A more interesting and relevant policy experiment is to imagine alternative taxes which generate the same revenue and thus generate the same income effect. For example, beginning with a simple life-cycle model with a proportional income tax, we could imagine replacing the income tax with a consumption tax, such as a broad-based GST. If the two

[8] Note that it is the taxation of *interest* income which generates this distortion to the after-tax real interest rate, not the existence of income tax in general. An income tax which applied only to labour income would not distort the return to saving (though it would clearly have other distortions).

taxes were designed to raise the same revenue, they would generate the same income effect. But the substitution effects would be very different. Indeed, since a consumption tax is levied on current and future consumption at the same rate, it does not distort the return to saving. In other words, the intertemporal relative price (after-tax interest rate) in a world with a consumption tax is exactly the same as in the world with no taxes at all. Thus, the act of replacing a proportional income tax with a broad-based consumption tax would generate a substitution effect which reduces current consumption and thus increases current saving.

B. RRSPs in the Simple Life-Cycle Model.

The conventional wisdom in the public finance literature is essentially that the introduction of RRSPs to an existing system which taxes interest income is equivalent to replacing the income tax with a consumption tax. This transformation of an income-tax system into a consumption-tax system is alleged to take place because the RRSP defers taxes from when the income is earned to when the income is spent. This conventional view therefore holds that the introduction of RRSPs successfully eliminates the distortion to the after-tax interest rate caused by the income tax—RRSPs therefore generate a substitution effect which unambiguously increases the amount of private saving.

This conventional view is not as simple as it first appears. The simple fact that RRSPs redistribute taxable income across periods is not sufficient to transform an income-tax system into a consumption-tax system. Put differently, the introduction of RRSPs into a system with an existing income tax does not necessarily eliminate the distortion caused by the taxation of interest income.

As Ragan (1994) has emphasized, the key point relates to the way in which RRSPs are introduced into the economy. If RRSPs are introduced as *the only instrument* for saving, then the conventional view is entirely correct. Since all saving is done through the RRSP, it follows that all income is taxed when the income is spent rather than when it is earned, and thus the income-tax system is effectively transformed into a consumption-tax system.[9] If RRSPs are

[9] There is one slight difference, however, which emphasizes that RRSPs really only permit the *deferral* (rather then the avoidance) of income taxation. With the current RRSP program, individuals are forced either to withdraw their funds by age 71, or roll their RRSP into a RRIF which pays an annuity. In either case, the individual is forced

instead introduced as *an additional instrument* for saving, then attention must be paid to the portfolio decision faced by individuals—they must decide how much to save inside their RRSP and how much to save outside the RRSP.

The policy experiment with real-world relevance is clearly the one in which RRSPs are introduced as an *additional* saving instrument—there have certainly been no suggestions that individuals be forced to do all of their saving inside of RRSPs. We thus confine the analysis to this case.

When individuals have access to two saving instruments, the individual faces an important portfolio decision. Saving outside of the RRSP earns a rate of return given by the after-tax real interest rate. Saving inside the RRSP is more complicated; funds inside the RRSP accumulate tax-free (being taxed only on withdrawal from the RRSP) and also earn the contributor a tax rebate in the year of the contribution. A wealth-maximizing individual, however, will contribute to each saving instrument so that the rate of return *for the marginal contribution* is equated across the two.

Since the (after-tax) real rate of return will be equated across the two saving instruments,[10] it follows that this single rate is the one which represents the relative price of current consumption in terms of future consumption. This is also then the single rate which determines the substitution effect (if any) generated by the introduction of RRSPs. Thus, the introduction of RRSPs will generate a substitution effect if and only if it changes this return to saving *at the margin*. Consider three situations.

• Proportional Income Taxes. Let $\tau(y)$ be the amount of taxes paid, where y is taxable income. Positive marginal tax rates imply $\tau'(y) > 0$. A proportional income-tax system therefore has $\tau'(y) > 0$ and $\tau''(y) = 0$. If r is the pre-tax real interest rate, then the after-tax interest rate is equal to $r_{AT} = [1 + r(1-\tau'(y))]$. With a flat-rate income tax, however, $\tau'(y)$ is a constant and is independent of the amount of saving done by the individual, inside or outside of the RRSP. The introduction of RRSPs thus generates no substitution effect at all.

---

to pay tax on these funds even when he might not be interested in spending them. In contrast, a pure consumption tat would never tax any funds until expenditure takes place.

[10] We address shortly the issue of corner solutions, when this marginal condition is violated.

In this flat-tax case, the marginal benefit of a one-dollar contribution to the RRSP always exceeds the marginal cost — RRSPs essentially represent an interest-free loan from the government to the individual who then has the incentive to contribute an unlimited amount to the RRSP. If there exists some upper limit to RRSP contributions, the individual contributes the maximum amount. But nothing will change the relative price of current consumption and thus there will be no substitution effect from the RRSP.

• Progressive Income Taxes. Suppose now that the income-tax system is progressive, and that this progressivity is achieved by having smoothly increasing marginal-tax rates. That is, $\tau'(y) > 0$ and $\tau''(y) > 0$. In this world, the after-tax return on saving outside the RRSP is equal to $r_{AT} = [1 + r(1-\tau'(y_F))]$, where y_F is the amount of taxable income in the future. Since r_{AT} is the after-tax interest rate linking the current period with the future, it is intuitive that the marginal income-tax rate *in the future* is central for determining r_{AT}.

Since the wealth-maximizing individual contributes to the two saving instruments to equate the marginal return, r_{AT} is the relevant relative price which determines the pattern of consumption and saving. Thus, RRSPs will generate a substitution effect if and only if the act of contributing to the RRSP brings about a change in r_{AT}. Clearly, this can happen only if current contributions to RRSPs change the level of taxable income in the future and thereby change the future marginal income-tax rate. And this is exactly what contributing to RRSPs in the current period does—it redistributes taxable income away from the current period and toward the future period. Thus, in a world which contains RRSPs (to which individuals contribute), y_F will be higher than it otherwise would be. This implies that RRSPs will raise the marginal income-tax rate in the future ($\tau'(y_F)$) and thus lower the after-tax interest rate. In a world with smoothly progressive income taxes, RRSPs therefore generate a *perverse* substitution effect—that is, they change the intertemporal relative price in such a way which *reduces* saving.[11]

• Asymmetric Treatment of Interest. The two previous cases implicitly assumed that interest was treated symmetrically by the tax system—that is, that interest income is taxed and interest expenses are tax deductible. A more realistic description of the Canadian tax system,

[11] If the tax system contains tax brackets, then $\tau''(y)=0$ within each bracket but $\tau''(y)>0$ across brackets. In this case, the perverse substitution effect depends on whether the RRSP contribution pushes the individual into a higher future tax bracket than otherwise would have applied.

however, would contain an asymmetry; interest income is taxable but interest expenses are not tax deductible. This asymmetry has important implications for the substitution effect of RRSPs.[12]

With this asymmetry, there is a kink in the individual's budget line. Individuals who are borrowers when they are young face an interest rate of r, whereas individuals who are lenders when they are young face an interest rate of $r(1-\tau'(y))$. The introduction of RRSPs is irrelevant for the young borrowers, but is important for the young lenders. For the young lenders, RRSPs produce a saving-inducing substitution effect, but only for those individuals who make less than the maximum allowable RRSP contribution—what might be called "interior" contributors. For individuals who are pushed up against the RRSP limit, any marginal saving for them must occur outside of the RRSP and thus the relevant interest rate for determining their consumption pattern is the (distorted) after-tax interest rate. For such "limit" contributors, therefore, there can be no saving-inducing substitution effect from RRSPs.[13]

C. The Effect of RRSPs on Private Saving.

The foregoing discussion emphasizes the role of the substitution effect in determining the influence of RRSPs on private saving. And, indeed, if the government introduced RRSPs into an income-tax system *while simultaneously imposing a tax to remove the generated income effects*, then the total effect of the RRSP on private saving would be given by its substitution effect.[14] However, for the broader policy question—that pertaining to the overall relationship between RRSPs and saving—it is probably more useful to think about the combined income and substitution effects on individuals' saving behaviour.

---

[12] In the United States, many interest expenses are not tax deductible, but mortgage interest is tax deductible. Given the importance of saving through home equity, the symmetric case probably applies more to the U.S. than to Canada.

[13] The conclusion is slightly altered in a world of asymmetric treatment of interest *and* smoothly progressive income taxes. The increase in the return to saving for a young lender must be offset by the increase in future marginal income-tax rates caused by the RRSP contribution. For a limit contributor, there will be a perverse (i.e. saving-reducing) substitution effect if future marginal income-tax rates are increased.

[14] This is the experiment which is implemented theoretically in Ragan (1994) and empirically in Section 4 below.

The income effect generated by RRSPs is relatively straightforward. Access to RRSPs makes the individual wealthier, for two reasons. First, by contributing to the RRSP in years of high marginal income-tax rates and withdrawing from the RRSP in years of lower marginal income-tax rates, the individual smoothes these rates across years. This reduces his lifetime income-tax obligations. Second, since the funds within the RRSP benefit from tax-free accumulation, this also implies an increase in wealth. If we assume that both current and future consumption are normal goods, then we expect the income effect from RRSPs to lead to higher current and future consumption. In what follows, we restrict our attention to the case where both current and future consumption are normal goods.

What is the total effect of RRSPs on private saving when we combine the income and substitution effects? This is perhaps best seen by considering a simple example. Consider an individual who has no access to RRSPs and is currently saving $5000 in some interest-bearing account. Now introduce a system of RRSPs. We begin by assuming that the individual transfers all $5000 from his saving account into the RRSP so that his total flow of saving is unchanged. Further assume that this individual faces a marginal income-tax rate of 40 percent. Since his taxable income has fallen by the amount of the RRSP contribution ($5000), he then receives a tax rebate equal to $2000.

It is helpful to focus on this $2000 increment to disposable income.[15] In particular, how does the individual spend this extra $2000? Consider two possibilities:

• Tax Rebate is Fully Saved. One extreme possibility is that the individual saves the full amount of the tax rebate. The individual's private saving thus increases by $2000 and the individual's current consumption is unchanged. In this case, current consumption is apparently not a normal good (since higher wealth does not increase current consumption), but future consumption is.

• Tax Rebate is Fully Consumed. Another extreme possibility is that the individual increases current consumption by the full amount of the tax rebate. The flow of private saving is therefore not affected by the RRSP but the flow of current consumption is. In this case, current consumption is a normal good, whereas future consumption is not.

[15] This $2000 increase in disposable income is not all new wealth; some of this will be required to pay taxes when the individual withdraws funds from the RRSP in the future.

It is tempting to view these possibilities as logical extremes which place bounds on the possible outcome. But there is nothing preventing the actual outcome from lying outside of these extremes. For example, it is conceivable that the increase in wealth (generated by the RRSP contribution) leads the individual to *reduce current consumption* and thus increase the flow of private saving by more than the $2000 tax rebate. Conversely, it is possible that the increase in wealth leads the individual to *reduce future consumption* and thus increase the flow of current consumption by more than the $2000 tax rebate. This requires the total flow of private saving to fall as a result of the RRSP.

Though both of these outcomes are logically possible, it is important to keep in mind what each of them requires in terms of the underlying income and substitution effects. In the first case, where current consumption actually falls, the substitution effect must reduce current consumption by more than enough to offset the income effect. In the second case, where current consumption rises by more than the tax rebate, the income effect is so strong that it dominates any substitution effect. Indeed, unless the substitution effect is very strong, this outcome is only possible if future consumption is an inferior good.

Now consider the three different types of individuals that we considered in the previous section: non contributors, "interior" contributors, and "limit" contributors. The combined income and substitution effects for these three types are:

- Non Contributors: No effect at all on private saving.

- Interior Contributors: The substitution effect for these individuals reduces current consumption.[16] The income effect—assuming consumption is normal in all periods—increases current consumption. The total effect on current consumption can thus go either way. If current consumption remains unchanged, then private saving rises by the full amount of the tax rebate. If current consumption rises (falls), then private saving rises by less (more) than the amount of the tax rebate.

[16] If future marginal tax rates rise significantly as a result of the RRSP contribution, then this substitution effect increases current consumption.

• Limit Contributors: The substitution effect for limit contributors either leaves current consumption unchanged or increases it. The income effect increases current consumption. The total effect therefore clearly increases current consumption. Since current consumption rises, private saving rises by less than the full amount of the the tax rebate.

To arrive at the total effect of RRSPs on private saving—aggregated over all individuals —we must combine the various income and substitution effects across these three types of individuals. It is clear that the *only* effect which reduces current consumption is the substitution effect for interior contributors. This consumption-reducing substitution effect must compete against the income effect for the same interior contributors and against both the income and substitution effects for limit contributors. Thus, the only way that aggregate private current consumption can fall when RRSPs are introduced is if the substitution effect for the interior contributors is strong enough to outweigh the combined income effects for all contributors. This may require a degree of intertemporal substitution in consumption on the part of these interior contributors that is implausible.

In summary, it seems reasonable to expect that RRSPs increase the flow of private saving, but also increase the flow of current private consumption. The two are simultaneoulsy possible because of the tax rebates earned by individuals contributing to RRSPs.

D. The Effect of RRSPs on Public Saving.

The tax rebate to the individual RRSP contributor, when looked at through the eyes of the government, appears as a reduction in income-tax revenue. For a given level of government expenditure, this reduction in income-tax revenues implies a reduction in government saving.

It is worthwhile noting that nothing literally implies that the level of government spending actually is constant when RRSPs are introduced; only that for analytical purposes it is useful to distinguish between changes in government saving which have their source in RRSP-generated changes in revenues from those which have their source in changes in government spending.

The effect of RRSPs on the level of government income-tax revenue has three components. First, there is the loss of tax revenue from the tax exemptions of the current flow of RRSP contributions. Second, there is the loss of tax revenue from the capital income earned within the RRSP. Finally, there is an increase in tax revenue from the taxation of current withdrawals from the RRSP.

Table 2 shows the magnitude of these tax expenditures for the federal government for 1991 and 1992. The top half of the table shows the tax expenditures for RRSPs; the bottom half shows the same for the employer-sponsored pension plans. For 1992, the federal government tax expenditures were $5.5 billion for RRSPs and $8.1 billion for RPPs. Note that these are only the *federal* tax expenditures; the associated numbers for the combined provincial governments would be similar in magnitude.

There are two points to note about the data on tax expenditures from Table 2, both of which introduce the possibility that Table 2 may not accurately reflect the *long-run* reduction in government income-tax revenue generated by RRSPs. The first relates to the maturity of the RRSP system; the second relates to the potential effects of RRSPs on the stock of income-earning assets. We address these in turn.

• Maturity of the System. The current system of RRSPs in Canada might be viewed as an immature system in the sense that the program is relatively young. The program started in 1957, but as is clear in Figure 2, RRSPs did not attract large flows of private saving until the early 1970s. Thus, relatively young people who began in the 1970s contributing to RRSPs for their retirement have probably not yet begun the process of withdrawing from the RRSPs. This is reflected in Table 2 by the much larger tax expenditures from the flow of current contributions than the taxes earned (negative tax expenditures) from the taxation of RRSP withdrawals.

What would Table 2 look like if the RRSP system were fully mature? That is, if the program had been in place with unchanged rules for long enough so that the *relative magnitudes* of the numbers in Table 2 were approximately constant over time. The answer depends on comparing the real interest rate (r), the rate of real per capita income growth (y), and the rate of population growth (n).

To see this, note that in a world in which an unchanging fraction of the population uses RRSPs, and in which the generosity of the program is unchanging, the annual flow of RRSP contributions will grow roughly at the rate n+y; that is, as individuals get richer and more numerous, there will be more RRSP contributions, and consequently more tax-revenue loss for the government. Similarly, a higher real interest rate means that the accumulation of RRSPs will be greater and hence the withdrawals will be greater. Thus, in an economy where r=y+n, the tax-revenue loss from contributions will approximately equal the tax-revenue gain from withdrawals. (Consider the special case in which r=y=n=0; in this case, we know that contributions exactly equal withdrawals, as long as the demographics are stable.)

To the extent that r=y+n as a long-run approximation, a mature system of RRSPs will therefore display a loss of tax revenues from RRSP contributions roughly equal to the gain in tax revenues from RRSP withdrawals. In such a mature system of RRSPs, only the tax expenditures associated with the untaxed capital income will be actively reducing the flow of government saving. In 1992, this amount was $2.8 billion for the federal government.

• The Stock of Assets. Our analysis of the loss in government tax revenue from RRSPs is so far only a "first-round" analysis in the sense that it looks at the implications for public saving coming only from the change in the current flow of tax revenues. Except for the discussion of the non-taxation of capital income earned within the RRSPs, the analysis has not addressed any possible effects that RRSPs might have on the stock of income-earning assets, such as the private capital stock. To the extent that RRSPs lead to a change in the private capital stock, such changes have an implication for the flow of future profits and thus the flow of future corporate income-tax revenue.

This point has been raised by Feldstein (1995) within the U.S. debate surrounding the effects of Individual Retirement Accounts (IRAs). Most of the focus in that debate has been on the effect of IRAs on the flow of private (as opposed to national) saving. Feldstein, however, argues that the increase in private saving generated by IRAs has the effect of increasing the private capital stock and thus increasing future government revenue from corporate income taxes. He thus argues that IRAs may well increase the flow of national saving in the United States.

Feldstein's argument, however, appears to rely on an unusual view of the "separation" of a country's capital market into private and public components.[17] Feldstein argues that the increase in private saving generated by IRAs is sufficient to increase the domestic capital stock, independent of what initially happens to the flow of government saving. This increase in the domestic capital stock in turn provides a larger corporate-tax base and thus results in increased government saving. A view of an integrated domestic capital market, in contrast, would suggest that if government saving initially falls by more than the increase in private saving, then the flow of national saving would fall. This reduction in national saving (under the Feldstein-Horioka view of limited international capital mobility) would then imply an eventual reduction in the domestic capital stock.[18]

E. Summing Up: Private Saving + Public Dissaving.

We have shown above the reasons to expect RRSPs to increase the flow of private saving. Moreover, unless the substitution effect faced by interior contributors is substantial, RRSPs will not only increase the flow of aggregate private saving but also increase the flow of aggregate private consumption. We have also shown the effect of RRSPs on government tax revenues; for a given level of government expenditure, RRSPs lead to a reduction in government revenues and thus a reduction in public saving.

Since private saving rises and public saving falls, it appears that the effect of RRSPs on national saving is ambiguous. But this ambiguity is only illusory. Recall from above that if RRSPs have the effect of increasing the flow of private consumption, then the increase in private saving is less than the individual's tax deduction. But this tax deduction is simply the size of the reduction in government saving. Thus, if RRSPs increase private consumption, then they must increase private saving by less than they decrease public saving—national saving falls.

---

[17] This differs from his view that capital has only limited international mobility.

[18] Ruggeri and Fougere (1995) have suggested that the closer integration of the corporate and personal income taxes in Canada than in the United States also weakens the applicability of Feldstein's argument to Canada.

Gale and Scholz (1994) provide some evidence in favour of this overall effect in the United States. They examine the effects on saving from increasing the annual contribution limits to IRAs from their 1983-85 levels. Similar to the discussion above, Gale and Scholz stress the importance of the fraction of the accompanying tax rebate that individuals save. They find that (1994, p.1235):

> *"...of the increased IRA contributions that would have resulted from increases in contribution limits, roughly 2 percent would represent net additions to national saving, if the accompanying tax cut were entirely saved. If one half of the tax cut were consumed in the first year, this estimate falls to -14 percent."*

Similar empirical studies for Canada—specifically addressing the effect of RRSPs on national saving—are only beginning to emerge. Section 4 presents the results of a recent study by Sargent (1995) which examines this important issue.

III. A Brief Look at Aggregate Data

In this section, we take a brief look at aggregate saving data. Observed patterns in various measures of Canadian aggregate saving have been discussed elsewhere, and a review of that work is not needed here. Jump and Wilson (1986) provide a general discussion about aggregate saving in Canada. Carroll and Summers (1987) discuss differences between Canadian and U.S. saving rates. Beach, Boadway and Bruce (1988) examine the relationship between taxation and saving.

A. Aggregate Saving in Canada, 1950-1994.

There are several components of aggregate saving in Canada. Personal saving is the saving of persons and unincorporated businesses. Corporate saving is the saving of private corporations and government enterprises. The sum of personal and corporate saving is private saving. Government saving is the saving of the consolidated government sector. National saving (sometimes called domestic saving) is then the sum of private saving and government saving.

Figure 3 shows the time-paths for personal, corporate and government saving in Canada from 1950 to 1994. Each time series is computed as gross saving as a share of Gross National Product (GNP). The path of national saving is shown in Figure 4.

Figure 3 shows the three components of national saving moving broadly together from 1950 to 1970, but taking quite different paths after the early 1970s. Specifically, the personal saving rate increased from a relatively constant average of 4 percent in the first twenty years of the sample to a high of over 12 percent in 1982 and then declined to its current value of about 5 percent. In contrast, the government saving rate averaged about 3 percent in the 1950-70 period and then began a dramatic downturn over the next 25 years, currently at around -4 percent of GNP. Only the corporate saving rate has shown no clear trend over the 45-year period; it averaged about 4 percent in the pre-1970 years and only slightly less in the later years. It clearly became more volatile in the post-1970 period, however, dropping by roughly 4 percentage points in the recessions of the early 1980s and the early 1990s.

Figure 4 shows the path of the national saving rate. Though the time path of national saving over the past half century is not smooth, there is a clear downward trend (albeit interrupted by a general increase from 1970 to 1980). From 1950 to 1955, the saving rate averaged about 12 percent; from 1960 to 1965, it averaged just under 10 percent; but from 1985 to 1995, it averaged only about 5 percent.[19]

There are three points which should be made about the measures of saving shown in Figures 3 and 4. The first relates to the general approach to measuring saving; the second relates to the distinction between gross and net saving; the third relates to inflation adjustment.

• Flows Versus Change-in-Stocks. The saving data shown in Figures 3 and 4 is drawn from the National Income and Expenditure Accounts and as such is based on the difference between the measured flow of income and the measured flow of expenditure.[20] It is well known, however, in Canada and elsewhere, that such flows-based measures of aggregate saving have the potential for serious problems (e.g., Dagenais 1992). An alternative measure of saving is to

[19] Though the precise measure of saving is different, the saving rates in Figure 4 are broadly consistent with those for Canada in Figure 1.

[20] The actual time series (levels, current dollars) come from CANSIM matrix #6705. Thet are then deflated by current dollar GNP.

use the national balance sheets to construct the change in the stock of assets from one year to the next. For example, unrealized capital gains on financial or housing assets clearly increase an individual's wealth—though they do not increase measured income in the national accounts—and thus can be expected to change the individual's consumption and saving behaviour. Thus, a one-dollar increase in the value of assets which *does not* lead to an increase in consumption would imply a one-dollar increase in the change-of-stock measure of saving.

In Section 4, where we examine household-level data on consumption and saving, we construct both the flow and change-of-stock measures for personal saving.

• Gross Versus Net Saving. In Figures 3 and 4 we use gross measures of saving deflated by gross measures of income. With this approach, any positive saving corresponds to foregone current consumption, even if that saving is being used to replace depreciated assets. By using a net concept of saving, in contrast, any saving which is used to replace depreciating capital stock (e.g., Capital Consumption Allowance) is subtracted from the gross measure to arrive at the net measure of saving. But this saving would still represent forgone current consumption.

Note also that the saving measures are shown as shares of GNP rather than GDP, reflecting the fact that GNP better represents the total amount of income accruing to domestic residents (as opposed to the total value of goods produced within the country).

• Inflation Adjustment. Inflation is well known to distort measures of saving. The basic reason (see Beach, Boadway and Bruce 1988) is that inflation erodes the real value of fixed-value assets such as bonds or cash balances, and this erosion of the real asset value is included in the measured flow of income. Even though there may be a one-for-one increase in nominal interest rates when inflation rises, measured nominal saving from the national accounts will overstate "real" saving by creditors, and will understate "real" saving by debtors. Thus, inflation tends to generate an overstatement of measured personal saving but an understatement of measured corporate and government saving.[21]

With this in mind, some of the large increase in the measured personal saving rate from 1970 through 1985 can be attributed to the significant inflation over this period. Similarly, some of the large decline in public saving over the same period is also due to inflation. An adjustment

[21] This is the familiar argument made by Robert Eisner (1989) that the national-accounts measures of government deficits should be corrected for inflation.

for inflation which yields the time series for "real" private and government saving would therefore reduce (but not eliminate: see Jump and Wilson 1986) the magnitude of the divergence in saving rates beginning in the early 1970s that is so clear in Figure 3. However, to the extent that the assets of the private sector are the liabilities of the government sector, then these two effects offset each other so that inflation does not distort the measure of national saving.[22]

B. Evidence About RRSPs and National Saving?

Carroll and Summers (1987) show that the early 1970s marked a dramatic departure in the behaviour of Canadian and U.S. private saving rates. Before the early 1970s, the two private saving rates moved closely together; after that time, the Canadian private saving rate embarked on a mild upward trend while the U.S. private saving rate trended downward. They argue that the main difference comes from the behaviour of the personal saving rates in the two countries, and suggest that the increasing generosity of the RRSP program in Canada at that time (see Table 1) may be an important explanation.

Their argument is suggestive. It is certainly reasonable to believe, mindful of the theoretical discussion in the previous section, that the greater generosity of RRSPs in Canada has been an important contributor to the rising personal saving rate in the past twenty years. After all, greater participation in RRSPs implies increases in disposable income, some of which will be spent and some of which will be saved.

Carroll and Summers (1987) did not examine the implications of RRSPs for government saving or for national saving. But a similarly suggestive argument can be made in this respect: The greater generosity of RRSPs has implied large reductions in government income-tax revenue and has thus been an important contributor to the sharp decline in the government saving rate over the past twenty years.

Furthermore, the theoretical arguments in the previous section give us some way to think about the likely overall effect on national saving. If individuals increase their saving by less than the reduction in government saving (the amount of the tax rebate)—that is, if individuals spend

[22] This offset will not be exact since Canada is an open economy. Some of the assets of the private sector are not the liabilities of Canadian government and, similarly, some of the liabilities of the private sector are not the assets of the Canadian private sector, being instead held abroad by foreigners.

any of their RRSP-generated tax rebates—then national saving can be expected to fall as RRSPs become more widespread. The increasing popularity of RRSPs may therefore be an important explanation for the decline in the *national* saving rate from 1970-94 that is so clear in Figure 4.

Though the broad patterns in the aggregate saving data are consistent with the view that RRSPs increase the flow of private saving but decrease the flow of government and national saving, Figures 3 and 4 taken alone clearly do not constitute solid evidence about the effect of RRSPs. Other factors could clearly be at work, and so there is an obvious need to explore the data in a more rigorous manner. Rather than looking at the aggregate data more closely, however, we instead review a recent study by Sargent (1995) which examines a large sample of households and estimates a formal model of RRSP saving and non-RRSP saving.

IV. An Empirical Model of Individual Saving

Sargent (1995) specifies an empirical model of household saving behaviour and estimates the model using a large and representative sample of Canadian households from 1992. Following the general approach of Gale and Scholz (1994) in the United States, Sargent specifically estimates the substitutability between RRSP saving and non-RRSP saving. If the two instruments are close substitutes, then a reduction in the generosity of RRSPs can be expected to increase non-RRSP saving substantially, with the result that total private saving changes by less than total RRSP saving. On the other hand, if RRSPs are viewed as poor substitutes for non-RRSP saving, then a reduction in the generosity of RRSPs may have little or no effect on non-RRSP saving; in this case, the change in total private saving will be approximately equal to the change in RRSP saving.

One advantage of this general framework is that once the basic estimation is complete, the model can be used to conduct real-world policy experiments (or simulations). Below, we discuss two such experiments, both of which were examined in the theoretical discussion in Section 2. First, we imagine the effects on individuals' saving behaviour of reducing the generosity of the RRSP program while holding constant their disposable income. This experiment essentially uncovers the substitution effect of RRSPs, and is precisely the policy experiment examined theoretically by Ragan (1994). The second experiment is to reduce the

generosity of the RRSP program while holding constant each individual's gross income. The results of this experiment reveal the overall effects (income plus substitution effects) of RRSPs.

The most striking results from Sargent's study can be summarized in the two following propositions.

> 1. If RRSPs are eliminated entirely, but personal taxes are lowered so that disposable income is held constant, total private saving will rise by about $86 per capita. In this case, the combination of policies is designed to maintain the level of government saving; the change in national saving is therefore exactly equal to the change in private saving, which in the aggregate equals about $840 million.

> 2. If RRSPs are eliminated entirely, but with no offsetting change in the personal income-tax rates, then national saving will increase by about $369 per capita, or $3.6 billion in aggregate. This overall increase comprises a small decline in private saving and a substantial increase in government saving.

In the remainder of this section, we present Sargent's (1995) study in some detail. The potential importance of the results requires a clear understanding of the basic empirical setting, upon which the discussion and interpretation in Section 5 can be based.

A. The Survey of Family Expenditures, 1992.

Sargent's study is based on the 1992 Survey of Family Expenditures (SFE). The survey is conducted every four years on a national basis and is designed to collect information about the flows of income and expenditure for private households. For 1992, the sample size is 9,492 households, and weights are provided so that the data can be made representative of the entire country. The unit of observation is the household (as opposed to the individual). For the purposes of estimating the formal model, Sargent uses a sub-sample of 5,279 households.

The primary advantage of using the SFE is that it contains information on both the household's flow of RRSP saving and non-RRSP saving. In contrast, the income-tax data includes data on RRSP contributions but not on the flow of non-RRSP saving, whereas the Survey of Consumer Finances has information on non-RRSP saving but not on RRSP contributions.

Gale and Scholz (1994) and Venti and Wise (1990) in the United States conduct their studies with data that do not contain direct observations on the flow of saving. Their use of the Survey of Consumer Finances forces them to assume a particular rate of return on both tax-sheltered saving and non-tax-sheltered saving in order to impute the annual flow of saving over several years from direct observations on the beginning-of-sample and end-of-sample asset stocks. The fact that the Canadian data in the SFE contain direct observations on the flow of both RRSP and non-RRSP saving may be viewed as a distinct advantage of Sargent's study relative to the similar studies in the United States.

• Two Measures of Saving. The SFE permits two different measures of household saving. The first is the net change in assets and liabilities, and is essentially equal to the change (during the year) in the household's net asset position. For example, if an individual reduced his holding of government bonds and used the proceeds to purchase a car, then this measure would show a decline in the individual's saving. Similarly, if an individual decreased the size of his mortgage by making a lump-sum payment to his mortgagor, then the increase in residential equity would appear as an increase in the individual's saving. We refer to this saving measure as the "change-in-stock" measure of individual saving. Note that *unrealized* capital gains on assets do not show up in measured income and also do not appear in this measure of household saving.

The second available measure of household saving—which might be called the "flow" measure of saving — is simply the observed difference between disposable income and consumption. Disposable income is defined to be pre-tax income minus personal income taxes minus mandatory contributions such as CPP/QPP and UI premiums.[23] Disposable income is therefore equal to take-home pay.

In principle, the two measures of saving are identical, although there are inevitable discrepancies due to measurement errors. It is perhaps surprising that the observed difference in the sample between the two measures of saving is only about $25 on average. Despite this almost trivial difference (on average), Sargent estimates the formal model and conducts all

[23] The household's contributions to employer-sponsored pension plans (RPPs) do not appear in the survey data. They are subtracted from gross income (as are UI and CPP premiums) and are similarly not included in the measure of the change in the value of assets.

policy experiments using both saving measures. This provides some check on the robustness of the results.

• RRSP Limits. Perhaps the biggest limitation with the data set (and this would be a problem with any available data set) is that the SFE contains no information on the RRSP limit faced by the household, or, more correctly, by the individuals in the household. The SFE does, however, report the earned income for the head of household and for the spouse.

Sargent uses the rules of the RRSP system together with some maintained assumptions to infer the RRSP limit for the household. First, he notes that in 1992 the RRSP limit was equal to 18 percent of earned income (up to a maximum of $11,500) minus the individual's Pension Adjustment. Sargent limits his study to those households with no more than two income earners (the head and the spouse). By doing so, he is able to estimate the limit for the two earners separately.

Sargent's next problem is in computing the appropriate Pension Adjustment for each individual. The SFE includes information on the *total household* contribution to employer-sponsored RPPs, but does not include contribution information for each individual separately, nor does it include information on the employer's contribution. Sargent makes two assumptions to deal with this problem. First, he assumes that the RPP contributions are split between the head and spouse according to each's share in total household *full-time labour income*. Second, he assumes that for each individual who contributes $1 to an employer-sponsored pension plan, the employer contributes $1.60. This number corresponds to the average ratio of employer-to-employee contributions in the 1992 National Income and Expenditure Accounts.

With these two assumptions, Sargent is then able to construct the Pension Adjustment for each individual. The resulting estimate in the household's RRSP limit is then simply the sum of the (up to) two individual RRSP limits.[24]

The final problem concerning the estimation of RRSP limits is that many households in the sample appear to contribute more than the upper limit, as constructed by Sargent. One possible reason for this observed over-contribution is that in 1992 individuals were permitted to

[24] Sargent uses individuals' 1992 income to construct their 1992 RRSP limit, but the actual 1992 limit was based on earned income in the previous year. Given that the SFE does not contain information on income from the previous year, this problem is unavoidable.

carry over their unused RRSP entitlement from the previous year, and this carry-over allowance is not built into Sargent's estimated limit (and it is not possible to do so with data only from 1992). There are two alternatives for dealing with this problem. The first is to effectively assume that the data is incorrect and that individuals really did not contribute more to their RRSPs than the (econometrician's) estimated limit. This is the approach used by Gale and Scholz (1994). The second alternative is to assume that the data is correct *but that individuals really did not over-contribute to their RRSPs*. In this case, it is the estimated RRSP limit which is incorrect. This is the approach used by Sargent. Thus, if L* is the estimated limit according to the method described above, and S is the actual amount contributed to the RRSP, then Sargent sets the RRSP limit to be

$$\text{RRSP Limit} = \max\{L^*, S\}.$$

(It is interesting to note that when Sargent uses the Gale-and-Scholz assumption for estimating the RRSP limit, his results are unchanged in terms of direction but much stronger in magnitude —that is, RRSPs are estimated to have a much stronger negative effect on saving.)

Table 3 (Sargent's Table 1) summarizes the saving behaviour for all households in the sample. The table shows that only 30 percent of the households were net contributors to RRSPs in 1992 (their contributions exceeded their withdrawals). Of those, nearly 26 percent contributed below their limit and 4 percent contributed at their limit. Perhaps not surprisingly, contributors had much more disposable income than non-contributors ($50,352 compared to $29,748).

There are four different measures of saving shown in Table 3. In each case, the dollar figure is the average per capita amount of saving; the percentage figure is the average amount of saving expressed as a percentage of average disposable income, and is thus a saving rate.

The first two saving measures are the "flow" and "change-in-stock" measures of total household saving discussed above. There is very little difference between these measures of saving, for either contributors or non-contributors. Contributors on average saved about 13 percent of their disposable income in 1992, compared to -1 percent for non-contributors. Among all RRSP contributors, the limit contributors saved the most, at around 20 percent of their disposable income.

The next two measures of saving in Table 3 show non-RRSP saving and RRSP saving, respectively. Obviously, contributors save more through RRSPs than do non-contributors (some

of which withdrew from their RRSPs and so the average saving in this cell is negative). But contributors also save considerably more outside the RRSP than do non-contributors. (Sargent shows several tables and figures which illustrate how saving is related to factors such as gross income and age of head of household.)

The final row in table 3 shows the estimated RRSP limit. What is perhaps the most striking thing here is the extent of under-contribution by the interior contributors, who on average are only contributing about one-third of their limit. So, despite the dramatic increase in the popularity of RRSPs that is evident in Figure 2, it is clear that the current system has a very considerable amount of unused RRSP "room".

B. The Formal Econometric Model.

Sargent's empirical specification is based on that used by Gale and Scholz (1994) in their study of U.S. saving. The basic approach is to model an individual's "desired" saving within the RRSP as well as their actual saving outside of the RRSP. An individual's actual RRSP saving will then equal their desired amount, unless that exceeds the RRSP limit, in which case their actual RRSP saving equals the limit. Similarly, if desired RRSP saving is negative, then actual RRSP saving is zero. Thus, RRSP saving is a "censored" variable, and this requires a (two-sided) Tobit estimation method.

Also important is the possibility of substitutability between RRSP saving and non-RRSP saving. In particular, if an individual is unable to contribute their desired amount to their RRSP (because it would put them above their limit), then this amount of "frustrated" RRSP saving may lead them to save more outside their RRSP than they otherwise would.

This basic setting is easy to represent formally. If S_R^* is the *desired* amount of RRSP saving, then

$$S_R^* = \mathbf{X}\beta + \upsilon$$

where **X** represents variables describing the individual, such as age, income, occupation, etc. The *actual* level of RRSP saving, S_R, is then assumed to be

$$= 0 \qquad\qquad S_R^* \leqslant 0$$

$$S_R = S_R^* = \mathbf{X}\beta + \upsilon \qquad 0 < S_R^* < L$$

$$= L \qquad S_R^* \geqslant L$$

where L is the actual RRSP limit. Sargent then models the non-RRSP saving behaviour by assuming that contributors and non-contributors are fundamentally different. Non-RRSP saving, S_N, is given by

$$= \mathbf{X}\gamma_1 + \epsilon_1 \qquad S_R^* \leqslant 0$$

$$S_N = \mathbf{X}\gamma_2 + \epsilon_2 \qquad 0 < S_R^* < L$$

$$= \mathbf{X}\gamma_2 + \eta(S_R^* - L) + \epsilon_2 \qquad S_R^* \geqslant L$$

where γ_1 describes the effect of **X** on S_N for non-RRSP-contributors and γ_2 describes the effect of **X** on S_N for RRSP contributors. $(S_R^* - L)$ is the amount of frustrated RRSP saving, and η measures the extent to which RRSP saving is substitutable with non-RRSP saving. η is itself estimated as a function of the **X** variables, $\eta = \mathbf{X}\bullet\delta$.

The parameter η is of central importance in this empirical model, and it is closely related to the issue of whether RRSP saving is primarily "new" saving or is just a movement of existing saving from non-sheltered accounts into the RRSP. If the estimated value of η is 1, then every dollar of frustrated RRSP saving gets saved outside the RRSP; in this case, a forced reduction in RRSP contributions (through a binding reduction in the limit) would lead to an equal increase in non-RRSP saving, with no overall change in private saving. Thus, $\eta = 1$ represents the case where RRSP saving is not new saving. At the other extreme, if the estimated value of η is 0, then a forced reduction in RRSP contributions would lead to no increase in non-RRSP saving and thus a reduction in overall saving. Thus, $\eta=0$ represents the case where RRSP saving is entirely new saving.

Sargent also notes that it is possible for the estimated value of η to exceed 1. This is a situation where one dollar of frustrated RRSP saving results in more than one dollar of non-RRSP saving. This might occur, for example, if the individual is a "target saver", whose goal is to have a specific level of resources available at some point in the future. In this case, since

RRSPs accumulate tax-free, the individual needs more than one dollar outside the RRSP for every one-dollar reduction in RRSPs, and thus η exceeds one.

C. Estimated Value of η.

Sargent's Table 10 (see his appendix) contains the estimation results. Since there are many variables included in **X**, it would be easy to spend considerable time interpreting the various parameter estimates in detail. Rather than doing this, however, we present here only a discussion of the estimated value of the central parameter η (for four separate cases). We then examine the policy simulations.

Table 4 shows the estimated value of η for four different cases. η is an estimated function of the **X** variables and so each individual (differing by X) has a separate estimated value of η. The values in Table 4 therefore reflect the average of the estimated η over all individuals. Furthermore, note that η is a behavioural parameter *only for limit contributors*, yet Table 4 shows the estimated value of η for all contributors. The interpretation of η for interior contributors is therefore the value of η that interior contributors would have if they were to (through some policy change) become limit contributors. That is, we could imagine a reduction in the RRSP limits which would transform a current interior contributor into a limit contributor, and that individual would then have an estimated value of η.[25]

The four separate cases correspond to four versions of the estimated model:

Model 1: Disposable Income; Flow measure of saving

Model 2: Disposable Income; Change-of-Stock measure of saving

Model 3: Gross Income; Flow measure of saving

Model 4: Gross Income; Change-of-Stock measure of saving

Keep in mind that these four models differ only with respect to the way that two variables are measured. The two different measures of saving are used only to provide some check on the

[25] We could also imagine the estimated value of η for current non-contributors, but this would only be relevant if we imagined a policy change which transforms a current non-contributor into a limit contributor. Such a policy experiment is not considered in Sargent's study, and so we ignore that possibility here.

robustness of the results. The two different measures of income are used because each corresponds to a distinct policy simulation. We discuss this more fully below.

The important point in Table 4 is the general pattern of η across interior contributors and limit contributors. First, note that the broad pattern is the same across all four models: η is higher for interior contributors than for limit contributors. In all cases, the value of η for limit contributors is well above zero, suggesting that there is considerable substitution between RRSP saving and non-RRSP saving.

It is significant that the estimated η is larger for interior contributors than for limit contributors. Imagine a case where the generosity of the RRSP program is substantially reduced. This has the effect of transforming some current interior contributors into limit contributors, who therefore experience some frustrated RRSP saving. The question then becomes: By how much do these frustrated limit contributors decide to increase their non-RRSP saving? Since the average value of η over these former interior contributors is greater than one, it is clear that for every dollar they are unable to put into the RRSP they will increase their non-RRSP saving by more than one dollar, thus raising total private saving. Thus, the estimated values of η in excess of one in Table 4 are central to understanding the overall results that reducing the generosity of RRSPs can be expected to increase overall saving in the economy. We turn now to the policy simulations which show these results.

D. Policy Simulations.

Each of the four versions of the estimated model can be used for the basis of several policy experiments. This naturally translates into many policy experiments—too many to keep the central message clear. Fortunately, the number of results presented can be immediately cut in half by noting that the estimation (and thus simulation) results are very similar for the two different measures of saving. We therefore present here only the results based on the "change-of-stock" measure of saving; Sargent presents the full set of results in his paper.

This leaves us with two versions of the estimated model; one based on disposable income and one based on gross income. These two models lend themselves to distinct policy experiments, both of which have interest in their own right. We examine each in turn.

• Gross Income Held Constant. Imagine reducing the generosity of the RRSP program (by reducing the RRSP limit) while at the same time making no offsetting changes in the personal income-tax system. Thus, reducing RRSPs implies an increase in the taxes paid by those who were originally contributing, and thus a reduction in disposable income. But if labour supply is assumed to be unaffected by this policy change, then there would be no change in gross income. This is the first policy experiment considered by Sargent. The simulations are based on the version of the model in which gross income is used as the measure of each individual's income.

Such a policy changes individuals' overall saving behaviour, and leads them to adjust their mix between RRSP and non-RRSP saving. These changes embody both the substitution effect and the income effect associated with the policy. The results are shown in Part A of Table 5. The first row corresponds to the simulated level of saving in the base case—with the actual 1992 RRSP program in place. The next three rows correspond to different policy experiments, each of which decreases the generosity of RRSPs. Sargent considers experiments that reduce RRSP limits by $1000, reduce RRSP limits by 50%, and eliminate RRSPs entirely.[26] The value shown in each column is the simulated dollar amount of saving under the policy experiment. All values are in per capita terms.

Consider the extreme policy experiment in which RRSPs are eliminated entirely, denoted Experiment 3. Private saving (the sum of RRSP and non-RRSP saving) is simulated to fall from $1,894 per capita in the base case to $1,835. Sargent then computes the associated change in government income-tax revenue by taking the reduction in RRSP contributions and applying the combined federal and provincial marginal income-tax rate. Under the assumption that the level of government spending is held constant, this increase in tax revenue implies a rise in government saving. Thus, the interpretation of the numbers in the second column is the change in government saving (in per capita terms) from the base case, where we normalize the level of government saving in the base case to be equal to $0. Thus, the value of $428 in the second column for Experiment 3 shows the total per capita tax expenditures (revenue loss) associated with the 1992 RRSP program.

[26] Sargent also considers an experiment in which RRSP limits are increased by $1000. The results are equal in magnitude but opposite in sign to the experiment in which limits are reduced by $1000.

The final two columns then show the effects on national saving (private plus government). Under the base case, private saving per capita is $1,894 and government saving is normalized to equal $0. With the elimination of RRSPs, private saving falls slightly to $1,835 per capita but government saving rises by $428 per capita. Thus, national saving rises by $369 per capita.

Note that "per capita" in this study means per household, of which there were approximately 9.8 million in 1992. Thus, the decline in private saving in this policy experiment is equal to $580 million whereas the increase in government saving is equal to $4.19 billion. The overall increase in national saving from this policy experiment is therefore equal to $3.6 billion.

One warning in terms of interpreting Table 5: Note that the aggregate value of tax expenditures from RRSPs in Table 5 ($4.19 billion) is considerably less than the actual amount from Table 2 ($5.5 billion). This is so for two reasons. First, since the SFE has data on the flow of contributions into RRSPs but not on the stock of RRSPs, Sargent's simulations cannot capture the change in government tax revenue from the taxation of already-sheltered funds (second row in Table 2). Thus, Sargent's simulations *understate* the increase in government saving which would result from the elimination of RRSPs. The second reason is more fundamental: the values in Table 5 corresponding to the Base Case are simulated values and thus due to the non-linear nature of the Tobit model do not equal the actual sample means. Table 5 should therefore not be used to examine the *level* of saving—only the *change* in saving resulting from various policy experiments.

• Disposable Income Held Constant. Now imagine quite a different policy experiment. The generosity of the RRSP program is reduced (again by lowering limits) but simultaneously the government adjusts the personal income-tax system in such as way that each individual's disposable income is held constant. Thus, since the reduction in RRSP generosity implies an increase in taxes for all former contributors, this policy involves a simultaneous tax reduction for each individual. Since disposable income is held constant by construction, any observed effects on private saving come only from the substitution effect of RRSPs. Furthermore, since government tax revenue is held constant by construction, any change in private saving implies exactly the same change (in dollar amount) for national saving.

Part B of Table 5 shows the results for this set of experiments. Note that there is no column for government saving since it is unchanged by construction. Instead, the table shows the break-down between RRSP and non-RRSP saving.

Consider again the extreme policy experiment in which RRSPs are entirely eliminated, Experiment 3. Obviously there is a reduction in RRSP saving from $1,020 per capita in the base case to $0. Non-RRSP saving rises from $875 per capita in the base case to $1,981 per capita. Note that non-RRSP saving increases by more than the reduction in RRSP contributions—this is the manifestation of the estimated value of η (averaged across all contributors) being greater than one. Total private saving therefore rises by $86 per capita. Since government saving is held constant in this experiment, the elimination of RRSPs increases national saving by $86 per capita, which implies an increase in national saving of $840 million.

V. Discussion

Section 2 makes the point that when individuals are given access to saving instruments such as RRSPs, it is important to think about the individual's decision regarding how much saving to do inside the RRSP and how much to do outside the RRSP. Only by modelling this choice is it possible to identify whether RRSPs change the rate of return at the margin, and thus whether they generate a substitution effect. As argued in Section 2, only the interior contributors in a simple life-cycle model would experience a substitution effect; limit contributors by definition are pushed up against a corner and thus RRSPs cannot generate a substitution effect for them. As Gravelle (1991) argues, however, the substitution effect is not the whole story. Any income effects associated with RRSPs must also be considered, and it is very likely that the income effects will decrease rather than increase saving.

The empirical exercises by Gale and Scholz (1994) in the United States and by Sargent (1995) in Canada show this same concern for understanding the individual's choice between competing saving instruments. These empirical models permit RRSP limits to be incorporated into the estimation process (hence leading to censored variables) and also permit the degree of substitution between RRSP saving and non-RRSP saving to be estimated. Moreover, once the models are estimated, they can then be used to conduct real-world policy simulations.

In what follows, we examine the consistency of Sargent's simulation results with the theoretical predictions from Section 2, the issue of Ricardian Equivalence, and the issue of what policy lessons might be gleaned from the entire exercise.

A. Consistency with Theoretical Predictions.

Sargent presents two extreme policy experiments which likely place bounds on what real-world policy makers might consider.

• Gross Income Held Constant. Sargent's first policy experiment is to simply reduce the generosity of RRSPs without any offsetting change in the personal income-tax system. Since reducing RRSPs implies reducing tax exemptions for contributors, this policy amounts to an increase in taxes for current contributors. The effects of this policy will therefore combine the income and substitution effects generated by RRSPs.

Consider the rather drastic policy of eliminating RRSPs entirely. What would the theoretical discussion of Section 2 predict to be the effects of such a policy? The substantial decrease in wealth for all current contributors suggests a decline in both current and future consumption. The substitution effect for current interior contributors is to consume more currently; all others have a substitution effect that is either zero or in the other direction. Summing over all individuals, the aggregate effect is very likely a reduction in current consumption (and a reduction in current saving, as there is also a fall in disposable income). The elimination of RRSPs provides a windfall to government tax revenues, however, equal to the value of current tax expenditures from RRSPs. The overall predicted effect is an increase in national saving—a rise in government saving which exceeds the decline in private saving.

Part A of Table 5 shows that this is exactly what happens in Sargent's policy experiments. For Experiment 3, private saving falls very slightly—by a little over 3 percent from the base case. This small reduction in private saving reflects the estimated value of η which is close to one; as RRSP limits fall to zero and all former contributors become frustrated RRSP savers, they choose instead to save roughly-equal amounts in non-RRSP saving. The former interior contributors (whose η exceeds one) save more now than before, but the former limit contributors (whose η is less than one) save less now than before. The net effect is less private saving, but only slightly. Government saving of course increases, so that the net effect

on national saving is an increase of $369 per capita. Summed over 9.8 million households, this implies an increase in national saving of $3.6 billion.

• Disposable Income Held Constant. At the other extreme, Sargent imagines a policy in which the generosity of RRSPs are reduced but the personal income-tax system is simultaneously changed so that each individual's disposable income is held constant. Though this is clearly an impossible-to-implement policy (getting the household-specific tax changes just right is not practically feasible), it is conceptually very useful because this policy by construction eliminates the income effect of the change in RRSP generosity. Thus, the simulated change in private saving reflects the pure substitution effect of RRSPs. Furthermore, this policy experiment holds public saving constant and thus the change in national saving is given by the change in private saving.

What would the theoretical apparatus in Section 2 predict if RRSPs were eliminated under these conditions? Since the income effects of the policy are eliminated by construction, we need only worry about the substitution effects. For non-contributors, this substitution effect is zero; for limit contributors, the substitution effect either has no effect or increases current saving; for interior contributors, the substitution effect implies more current consumption, and thus less saving. Summed (and averaged) over all individuals, therefore, the effect on private saving is likely to be very small, and could be positive or negative.

Part B of Table 5 shows that the substitution effect increases saving but only slightly. With the elimination of RRSPs, per capita RRSP contributions fall from $1,020 to zero. Non-RRSP saving rises from $875 to $1,981, an increase of $1,106. Again, the fact that non-RRSP saving rises by only slightly more than the fall in RRSP saving reflects the estimated value of η close to (but above) one. The former interior contributors increase their total saving as a result of the policy whereas the former limit contributors decrease their total saving. But there are many more interior than limit contributors and so the net effect is a slight increase in total private saving, $86 per capita. This implies an increase in national saving of $840 million.

B. The Role of Ricardian Equivalence.

Whenever a proposed policy has the effect of changing current tax revenue while maintaining unchanged the path of government spending, the issue of Ricardian Equivalence must be

addressed. The Ricardian Equivalence proposition, revived twenty years ago by Barro (1974) and actively discussed ever since, is that changes in the *timing* of taxation do not affect the consumption decisions of private agents since those agents can borrow or lend in the capital market as necessary to offset changes in the government deficit. Specifically, an increase in taxes which raises the flow of government saving, but which must be matched in present value by future tax reductions, will lead Ricardian consumers to reduce their private saving dollar-for-dollar and thus keep national saving unchanged.

In terms of the RRSP policy experiments examined here, note that Ricardian Equivalence pertains only to the income effect of changes in the RRSP program; it is only the income effect which changes the amount of government tax revenue and thus opens the door for questions concerning Ricardian Equivalence. This immediately implies that the policy simulation in which disposable income is held constant (by keeping total tax revenues constant) is not subject to any criticisms regarding Ricardian Equivalence.

But the gross-income-held-constant policy experiment leads to a substantial increase in government saving, and thus the issue of Ricardian Equivalence must be addressed. To examine this issue, consider what might be called the "Ricardian offset"—the degree to which private individuals are Ricardian in the sense that they offset tax-financed changes in the flow of government saving. A Ricardian offset of 0% means that agents are completely non-Ricardian; an offset of 100% means that they are purely Ricardian.

Table 6 shows the saving effects generated by the complete elimination of RRSPs under various values for the Ricardian offset. Each column shows the change in saving from the base case in Part A of Table 5 (that is, the first row in Table 6 corresponds to the last row in Part A of Table 5). In this case, with non-Ricardian consumers, private saving falls by $59 per capita and government saving rises by $428 per capita; national saving therefore rises by $369 per capita. As the Ricardian offset rises toward 100%, however, the decline in private saving increases in magnitude to offset more of the increase in government saving; individuals view the current rise in government saving as being matched in the future with equivalent tax reductions, and this leads individuals to lower their current flow of saving. Thus, as the Ricardian offset increases toward 100%, the increase in national saving from this policy declines. If individuals are purely

Ricardian, then they completely offset the increase in government saving, and the overall effect is to reduce national saving by $59 per capita.[27]

Perhaps the most surprising finding in Table 6 is the size of the Ricardian offset required in order to reverse the basic result that eliminating RRSPs (holding gross income constant) generates an increase in national saving. Even with a Ricardian offset of 80%, so that consumers are almost purely Ricardian, the overall decline in private saving is not as large as the increase in government saving; the overall effect on national saving in this case is still an increase equal to $27 per capita, or about $260 million in aggregate. Though the "correct" size of the Ricardian offset in Canada is not something economists yet agree on, most economists would probably believe that it is significantly greater than zero but also significantly less than 100.[28]

In summary, unless one takes an extreme view of the size of the Ricardian offset, Sargent's policy experiments suggest that the elimination of RRSPs, holding gross income constant, will increase the flow of national saving. A 50% Ricardian offset (to choose the midpoint of the range) implies an increase in national saving of $155 per capita, or $1.5 billion in aggregate. It is worth repeating here that the other policy experiment, in which disposable income is held constant, does not suffer from the analytical problems of Ricardian Equivalence; the predicted effect in that experiment is that national saving rises by $86 per capita or $840 million in aggregate.

C. Policy Lessons.

This paper began with a discussion of the importance of saving to the overall economy, and the point was made that in a closed economy saving was obviously necessary to finance domestic investment. Though this is not in principle true in an open economy with perfect capital mobility, there may be good reasons to believe that even open economies behave as if they are closed, in the sense that a significant share of any increase in domestic saving tends to remain in

[27] The fact that there is a non-zero effect on national saving even in the case of purely Ricardian consumers highlights the fact that Ricardian Equivalence relates only to the income effect; there is still a substitution effect which alters individuals' consumption and thus national saving, even when they completely offset changes in the flow of government saving.

[28] See David Johnson (1994) for a discussion of Ricardian Equivalence in Canada. See also the comments on Johnson's paper by Robert Lucas.

the country to finance domestic investment. Thus, policies which successfully promote domestic saving may be effective in increasing domestic investment and, through that channel, stimulating economic growth.

Saving is closely related to taxation. A conventional view of income taxes is that the taxation of interest income distorts the consumption-saving decision by lowering the after-tax real return to saving. Under this view, the income tax is a distortion which leads the private economy to save less than what is socially desirable (or Pareto efficient), and thus may be responsible for domestic investment being less than what it would be with a non-distortionary tax. One motivation for the introduction of tax-assisted saving plans (like RRSPs and RPPs), then, is to eliminate this distortion in the personal income-tax system and thus bring about an increase in national saving.

There are two central theoretical messages in this paper which help us to understand whether the introduction of RRSPs is likely to increase the flow of national saving.

• Substitution and Income Effects. First, though income taxes may indeed distort individuals' consumption-saving decisions, RRSPs may be unsuccessful in bringing about a transformation of the income-tax system into a consumption-tax system (which does not distort the intertemporal relative price). Here, we must think about the importance of limit contributors and about the possibility of perverse substitution effects generated by the progressive nature of the income-tax structure. But it is also important to think about the income effects of RRSPs; the substitution effect is important for understanding the extent to which RRSPs eliminate the existing distortion caused by income taxation, but the income effect is perhaps more important for understanding the effect on the flow of private saving. The arguments in Section 2 suggest that the average substitution effect over all individuals in the economy is likely to be very small. The income effect, in contrast, may be quite significant.

• Government Saving. The second central message is that to the extent that we care about the flow of national saving—and it is certainly *national* saving which affects the flow of domestic investment if capital is imperfectly mobile across international boundaries—then the analysis of RRSPs cannot stop with the behaviour of the individual. It is equally important to think about the effects of RRSPs on the flow of government saving. Indeed, the effects of RRSPs on government saving is exactly the flip side of the tax rebate received by individual

contributors, which gives rise to their income effect. The arguments in Section 2 suggest that RRSPs generate a significant reduction in government saving.

This theoretical analysis—buttressed by the simulations of Sargent's empirical model of individual saving—suggests that while RRSPs may indeed be an effective means of increasing the flow of *private* saving, they are unlikely to increase the flow of *national* saving. On the contrary, since RRSPs increase the wealth of individual contributors, those contributors are likely to increase their current consumption, resulting in an increase in private saving which is less than the reduction in the flow of government saving. Thus, not only are RRSPs an ineffective mechanism for increasing the flow of national saving, they are very likely responsible for decreasing it.

VI. Final Remarks

This paper has examined a specific and narrow issue—the effect of RRSPs on the flow of private and national saving. It has done so in a positive (as opposed to normative) way. Section 2 discussed the probable effects of RRSPs on private consumption, private saving and government saving. Section 4 examined in detail a recent study by Sargent (1995) which estimates a formal model of individual saving. The policy simulations conducted by Sargent show results broadly consistent with the predictions from Section 2. Sargent's study was also an exercise in positive analysis.

The conclusion of this paper is that there are solid grounds for believing that a reduction in the generosity of the RRSP program would lead to a reduction in private saving but a larger increase in government saving. The flow of national saving, therefore, can be expected to rise as RRSP generosity is curtailed.

Note, however, that the conclusion "RRSPs reduce national saving" is very different from the conclusion "RRSPs are undesirable". The first does not necessarily imply the second. Indeed, the relationship between RRSPs and national saving may be considerably less central than other issues related to RRSPs, and it important that this study be placed in a broader context.

This paper closes with a brief discussion of what other issues deserve to be addressed before policy makers can rule on the general "desirability" of RRSPs.

• Why Promote National Saving? Suppose, for the sake of argument, that RRSPs are responsible for increasing the flow of national saving. Before concluding that RRSPs are desirable in this case, it is important to know why the promotion of national saving is a sensible policy goal. If national saving is too low because of income-tax-caused distortions, then RRSPs may be desirable on the grounds that they reduce existing distortions (though there are probably better ways of achieving this). But if the accepted wisdom is simply that national saving should be higher so that future output and consumption can be higher, then it is important to keep in mind that RRSPs do not offer a free lunch. Even if RRSPs increase national saving (which is doubtful), they do so by reducing the flow of current consumption, and it not clear on what grounds the government should be promoting future consumption at the expense of current consumption.

This suggests that more attention be paid to examining the general issue of: "Why do we care about the level of national saving, and under what situations should government be taking active policy in this area?"

• A Possible Externality. National saving is not the only issue related to RRSPs. Though it is difficult to find in the public record a clear and precise motivation behind a policy of tax-assisted saving, some people believe that RRSPs are desirable on the grounds that they increase the flow of private saving (independent of what they do to the flow of national saving). From the perspective of the general taxpayer, such an increase in somebody else's private saving may be rightly viewed as desirable since it reduces the likelihood of that individual being a drain on the public purse in the future. In this view, there is essentially a positive externality associated with private saving—the more saving done by any given individual, the less other individuals collectively will have to support them when they are old.

If this is the motivation behind RRSPs, then an important question to examine is: "Who is contributing to RRSPs, and without such contributions would these individuals face a significant probability of requiring old-age income support from the government?" If the finding is that only the relatively well-to-do are contributing to RRSPs (as Table 3 in fact suggests),

then RRSPs may be missing their mark. In any event, this is an issue which certainly deserves to be explored in considerably more detail.

• Income Distribution. Aside from the effect of RRSPs on private or national saving, policy makers should be interested in knowing whether RRSPs generate any unintended side effects and, if so, whether these side effects are generally desirable. One possibility discussed in Ragan (1996) is that RRSPs may be responsible for redistributing income in the "wrong" direction. This is so for two reasons. First, as Table 3 suggests, it is typically the higher-income households that use RRSPs and thus benefit from the RRSP-induced tax exemptions. Second, higher income households typically face higher marginal income-tax rates and thus benefit more from the tax-free accumulation because they avoid a higher rate of taxation. As Ragan (1996) shows, there are considerable redistributions taking place from non-contributors toward contributors, and from low-income contributors toward high-income contributors.

Though it is always dangerous to focus on the redistributive property of individual elements of the tax system (as opposed to the entire tax system as a package), it is difficult to believe that this amount of regressivity is required within one element of an overall tax system that strives for progressivity. The general issue of the distributional implications of RRSPs certainly deserves to be examined more fully.

One final comment. The existence of so many non-contributors and interior contributors (see Table 3) suggests that the simple life-cycle model of saving is missing something important. After all, if RRSPs are able to increase an individual's wealth, it is more than a little puzzling that so few people use them. Though the basic insight of the life-cycle model—essentially viewing individuals as consumption smoothers—is extremely powerful and helps us to understand many issues in the economics of consumption, saving and taxation, it still leaves some puzzles unexplained. In an effort to understand some of these remaining puzzles, we should feel free to "go beyond" the basic life-cycle model, as has recently been explored by Thaler (1990, 1994), Hubbard et al (1994) and others.

References

Barro, R. (1974), "Are Government Bonds Net Wealth?", *Journal of Political Economy*, 1095-1117.

Beach, C., R. Boadway and N. Bruce (1988), *Taxation and Savings in Canada*, Economic Council of Canada.

Bosworth, B. (1989), "There's No Simple Explanation for the Collapse in Saving", *Challenge*, 27-32.

Carroll, C. and L. Summers (1987), "Why Have Private Savings Rates in the United States and Canada Diverged?", *Journal of Monetary Economics*, 249-279.

Carroll, C. and D. Weil (1994), "Saving and Growth: A Reinterpretation", *Carnegie-Rochester Conference Series on Public Policy*, 133-192.

Dagenais, M. (1992), "Measuring Personal Savings, Consumption, and Disposable Income in Canada, *Canadian Journal of Economics*, 681-707.

De Gregario, J. (1993), "Savings, Growth and Capital Markets Imperfections: The Case of Borrowing Constraints", International Monetary Fund Working Paper 93/31.

Eisner, R. (1989), "Budget Deficits: Rhetoric and Reality", *Journal of Economic Perspectives*, 73-93.

Feldstein, M. and C. Horioka (1980), "Domestic Saving and International Capital Flows, *The Economic Journal*, 314-329.

Feldstein, M. (1995), "Too Little, Not Too Much", *The Economist*, June 24.

Feldstein, M. (1995), "The Effect of Tax-Based Saving Incentives on Government Revenue and National Saving", *Quarterly Journal of Economics*, 475-494.

Hubbard, G., J. Skinner and S. Zeldes (1994), "Expanding the Life-Cycle Model: Precautionary Saving and Public Policy", *American Economic Review* (May), 174-179.

Gale, W. and J. Scholz (1994), "IRAs and Household Saving", *American Economic Review*, 1233-1260.

Gravelle, J. (1991), "Do Individual Retirement Accounts Increase Savings?", *Journal of Economic Perspectives*, 133-148.

Horner, K. and S. Poddar (1992), "Pension Reform in Canada", International Institute of Public Finance, mimeo.

Johnson, D. (1994), "Ricardian Equivalence: Assessing the Evidence for Canada", in *Deficit Reduction: What Pain, What Gain?*, C.D.Howe Institute.

Jump, G. and T. Wilson (1986), "Savings in Canada: Retrospective and Prospective", (MacDonald Commission Reports).

Krugman, P. (1994), "The Myth of Asia's Miracle", *Foreign Affairs*, November/December 1994.

Levine, R. and D. Renelt (1992), "A Sensitivity Analysis of Cross-Country Growth Regressions", *American Economic Review*, 942-963.

Nordhaus, W. (1989), "What's Wrong with a Declining National Saving Rate?", *Challenge*, 22-26.

Ragan, C. (1994), "Progressive Income Taxes and the Substitution Effect of RRSPs", *Canadian Journal of Economics*, 43-57.

Ragan, C. (1996), "A Case for Abolishing RRSPs", C.D.Howe *Social Policy Challenge*, forthcoming.

Ruggeri, G. and M. Fougere (1995), "The Effect of Tax-Based Savings Incentives on Government Revenue and National Saving: Comment", Department of Finance, mimeo.

Sargent, T. (1995), "RRSP and Non-RRSP Saving: An Empirical Analysis", Department of Finance, mimeo.

Solow, R. (1956), "A Contribution to the Theory of Economic Growth", *Quarterly Journal of Economics*, 65-94.

Thaler, R. (1990), "Saving, Fungibility, and Mental Accounts", *Journal of Economic Perspectives*, 193-205.

Thaler, R. (1994), "Psychology and Saving Policies", *American Economic Review* (May), 186-192.

Venti, S. and D. Wise (1990), "Have IRAs Increased U.S. Saving?: Evidence from Consumer Expenditure Surveys", *Quarterly Journal of Economics*, 661-698.

Young, A. (1992), "A Tale of Two Cities: Factor Accumulation and Technical Change in Hong Kong and Singapore", *NBER Macroeconomics Annual* (ed. O.J. Blanchard and S.Fischer).

Table 1

RRSP Contribution Limits, 1957-1994

Year of Policy Change	Individuals without RPPs	Individuals with RPPs
1957	10% of earned income to a maximum of $2500	10% of earned income to a maximum of $1500
1972	20% of earned income to a maximum of $4000	20% of earned income to a maximum of $2500
1976	20% of earned income to a maximum of $5500	20% of earned income to a maximum of $2500
1986	20% of earned income to a maximum of $7500	20% of earned income to a maximum of $3500
1991	18% of earned income to a maximum of $11500	18% of earned income to a maximum of $11500-PA
1993	18% of earned income to a maximum of $12500	18% of earned income to a maximum of $12500-PA
1994	18% of earned income to a maximum of $13500	18% of earned income to a maximum of $13500-PA

PA is the individual's Pension Adjustment, which reflects benefits from an employer-sponsored RPP. For individuals with a defined-contribution plan, PA is the sum of the employer and employee contributions. For a defined-benefit plan, an imputation is made, based on the details of the plan, to compute the PA. See Horner and Poddar (1992).

Table 2

Federal Personal Income-Tax Expenditures from RRSPs and RPPs

	1991	1992
	(billions of dollars)	
Registered Retirement Saving Plans		
Deductions for Contributions	3.310	3.685
Non-taxation of investment income	2.980	2.755
Taxation of withdrawals	-0.925	-0.940
Total Revenue Loss from RRSPs	5.365	5.500
Registered Pension Plans		
Deductions for Contributions	4.460	4.990
Non-taxation of investment income	8.950	7.690
Taxation of withdrawals	-4.030	-4.580
Total Revenue Loss from RPPs	9.380	8.100

Source: *Government of Canada Tax Expenditures*, December 1994.

Table 3

Saving Behaviour for All Households in SFE Sample

	Household Characteristic				
	All Households	Non Net Contributors	Net Contributors All	Interior	Limit
% of all households	100%	70%	30%	25.7%	4.3%
Disposable Income	$35,921	$29,748	$50,352	$50,114	$51,773
Flow Saving					
Average	$1,667	-$417	$6,538	$6,019	$9,645
Rate	4.6%	-1.4%	13%	12%	18.6%
Change-in-Stock Saving					
Average	$1,669	-$366	$6,428	$5,742	$10,536
Rate	4.6%	-1.2%	12.8%	11.5%	20.4%
Non-RRSP Saving					
Average	$794	$41	$2,554	$2,538	$2,646
Rate	2.2%	0.1%	5.1%	5.1%	5.1%
RRSP Saving					
Average	$875	-$408	$3,874	$3,204	$7,890
Rate	2.4%	-1.4%	7.7%	6.4%	15.2%
RRSP Limit (Estimated)	$5,617	$3,973	$9,460	$9,723	$7,890

Source: Sargent (1995)

Table 4

Estimated values of η

	Model 1	Model 2	Model 3	Model 4
Interior	2.03	1.39	1.20	1.26
Limit	1.45	0.90	0.72	0.64

Model 1: Disposable Income; Flow measure of saving
Model 2: Disposable Income; Change-of-Stock measure of saving
Model 3: Gross Income; Flow measure of saving
Model 4: Gross Income; Change-of-Stock measure of saving

Source: Sargent (1995)

Table 5

Part A. Policy Experiments Holding Gross Income Constant

Environment	Private Saving	Government Saving	National Saving	ΔNational Saving
	(all values are per capita dollars)			
Base Case	$1,894	$0	$1,894	$0
Experiment 1 RRSP Limits Reduced by $1000	$1,889	$14	$1,903	$9
Experiment 2 RRSP Limits Reduced by Half	$1,882	$65	$1,947	$53
Experiment 3 RRSPs Eliminated Entirely	$1,835	$428	$2,263	$369

Part B. Policy Experiments Holding Disposable Income Constant

Environment	RRSP Saving	Non-RRSP Saving	Private Saving	ΔNational Saving
	(all values are per capita dollars)			
Base Case	$1,020	$875	$1,895	$0
Experiment 1 RRSP Limits Reduced by $1000	$ 983	$917	$1,900	$5
Experiment 2 RRSP Limits Reduced by Half	$ 859	$1,046	$1,904	$9
Experiment 3 RRSPs Eliminated Entirely	$ 0	$1,981	$1,981	$86

Source: Sargent (1995)

Table 6

Saving Effects When RRSPs are Eliminated:
The Role of Ricardian Equivalence

Ricardian Offset	ΔPrivate Saving	ΔGovernment Saving	ΔNational Saving
(changes in per capita dollars from the Base Case, Part A of Table 5)			
Non-Ricardian Consumers			
0 %	-$59	$428	$369
20 %	-$145	$428	$283
40 %	-$230	$428	$198
60 %	-$316	$428	$112
80 %	-$401	$428	$27
Pure Ricardian Consumers			
100 %	-$487	$428	-$59

Figure 1. Cross-Country Saving and Growth Rates Over Three Decades

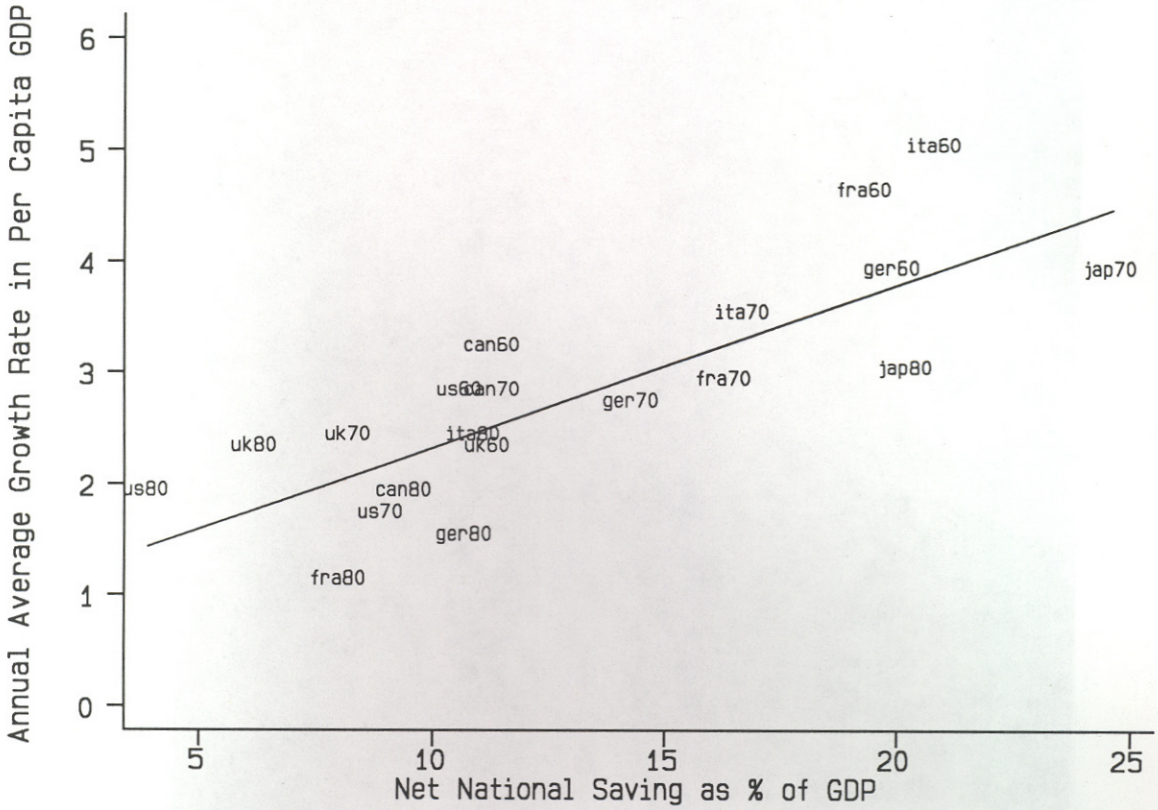

Figure 2. Annual Contributions to RRSPs and RPPs, 1960-1992

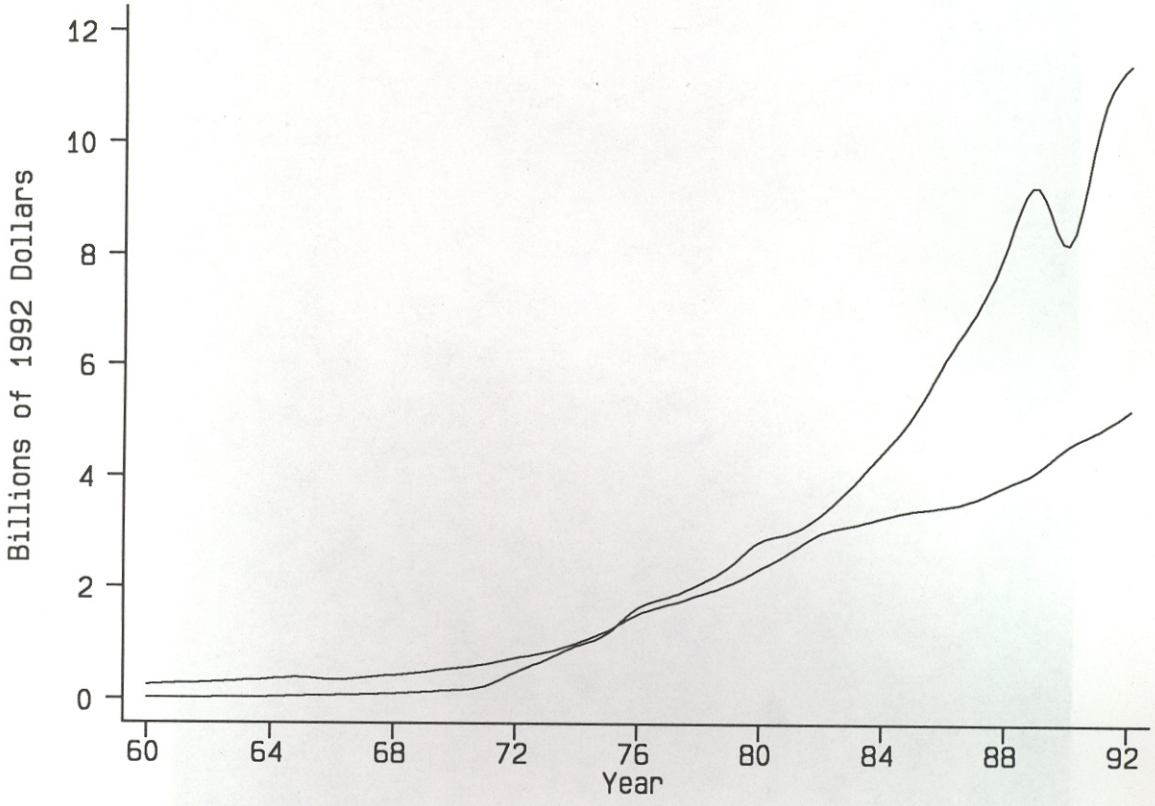

Figure 3. Personal, Corporate and Government saving, 1950-94

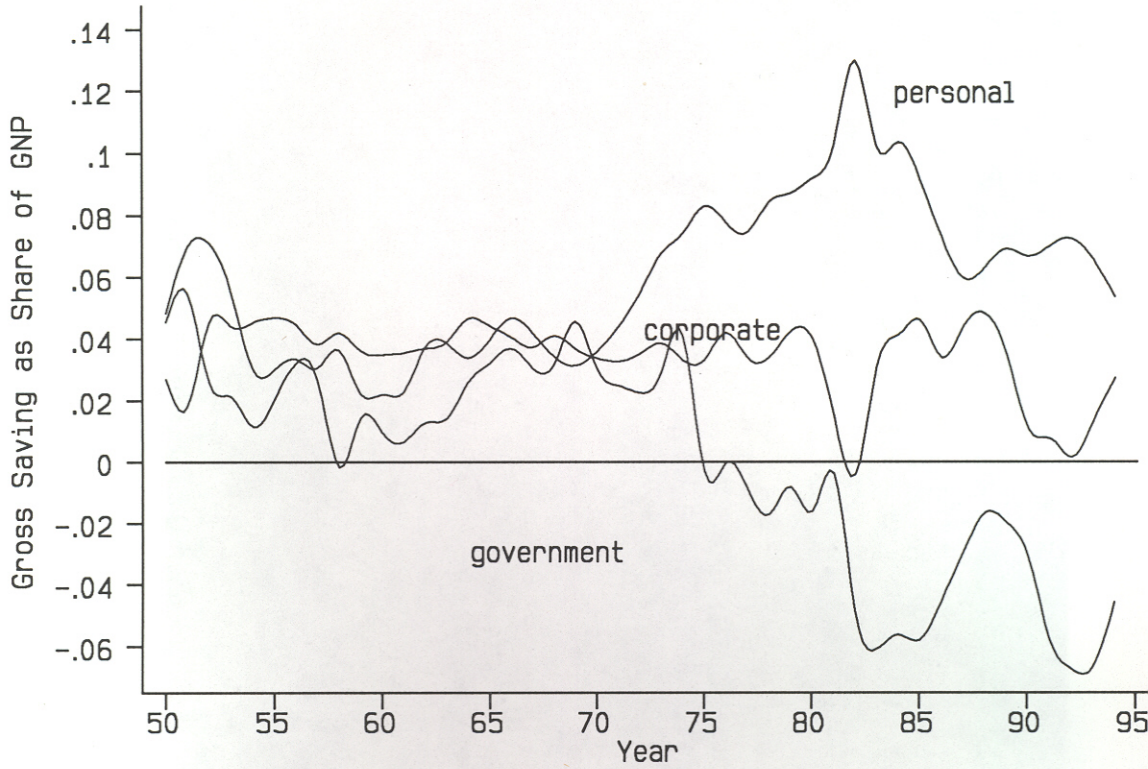

Figure 4. Canadian National Saving Rate, 1950-94