Like most countries,Canada has a public pension plan, the Canada Pension Plan. (I will denote the Canadian plan as CPP/QPP, since Québec residents have their own provincial plan, with a similar structure, instead of the federal plan.)

The CPP/QPP is certainly not the only public involvement in the fields of income security for old people, or saving. There are two major programmes providing income for old people, the Old Age Security ("OAS") programme, and the Guaranteed Income Supplement ("GIS"). Both are funded from general tax revenues. The OAS provides a pension for all people over the age of 65. A person's OAS benefits depend on her age, sex, and marital status, but not on her income. In contrast, the GIS provides extra benefits only to low–income old people. The amount of GIS benefits which a person receives depends on her income ; as her income rises, GIS benefits are "clawed back", until they reach 0.

As an inducement to saving for retirement, the federal and provincial governments offer tax incentives for people to save using Registered Pension Plans ("RPP"), and Registered Retirement Savings Plans ("RRSP"). RPPs are company pension plans, to which employers and employees both contribute. RRSPs are private plans, to which individuals may contribute. In each case, contributions to the plans are tax deductible in the year in which they are made (and for two more months in the subsequent year). Withdrawals from the plans are taxed but at a reduced rate if the income is withdrawn during retirement.

Why is CPP/QPP different from these other programmes? Unlike OAS/GIS, CPP/QPP is not supposed to be financed from general tax revenue, but from earmarked "contributions". That is, the CPP/QPP are supposed to be self–financing plans, in which payments to old people are not supposed to be taken from general tax revenues, but from money available within the CPP/QPP plan. Unlike RPP/RRSP, CPP/QPP is **compulsory**.

Most countries' public pension plans share these features of CPP/QPP. They tend to have special taxes, earmarked for pensions only, and the pension plans are supposed to be self–financing. They are also compulsory in most countries. The American Social Security plan, for example, shares both of these features.

How do these public pension plans differ from private pensions? As mentioned above, CPP/QPP (and U.S. Social Security) are compulsory. Many firms have private pension plans, which are compulsory for employees of these firms. So the difference here is not large. Of course, if one does not want to participate in some firm's generous (and expensive) pension plan, one can always quit the firm, and look for work with a firm with a less elaborate pension plan. In contrast, one can only escape CPP/QPP by leaving Canada, or by leaving the work force.

But there are other differences. Private pension plans tend to be classified as either **defined contribution**, or **defined benefit**. In recent years, defined contribution plans have become more popular.

In a defined benefit plan, workers (and often employers as well) pay into the plan during

the worker's career with the firm. The worker is then entitled to collect benefits when retired, according to some formula. The formula relates the benefits she will collect when retired to the number of years she worked for the firm, her salary with the firm, and other characteristics (such as sex or marital status). So in a defined benefit plan, workers are **guaranteed** a specific level of pension benefits, defined by this formula. Of course, the formula must be feasible. The assets of the company pension plan must be sufficient to provide the promised benefits. So defined benefit pension plans must be audited, in order to ensure that the plan does have adequate assets to meet its explicit obligations to future retired workers.

A defined contribution plan works somewhat differently. Each worker has her own pension account. The payments she makes into the plan while working, together with any payments made by the employer, go into her **own account**. When she retires, she is entitled to a company pension, based on the value of her own account. That is, she gets the return from the assets into which her own contributions (and those made by the employer on her behalf) have been invested.

So, among private pensions, a key difference between defined benefit and defined contribution plans is the allocation of risk. The return on financial assets is uncertain. In a defined contribution plan, the worker bears the risk : she gets a larger pension if the market (or the part of the market in which her contributions have been invested) does well, but a smaller pension if the market does poorly. In a defined benefit plan, her pension benefits are not tied to the market. So the employer bears the risk under a defined benefit plan : somehow the employer must make up the difference to the retired worker, if the market does badly, and the return on invested contributions is less than has been promised to workers.

By this taxonomy, the CPP/QPP is a defined benefit plan. That is, individual Canadian workers do not have their own CPP/QPP accounts. And the payments that retired people get from CPP/QPP do not depend on how high the return has been on assets which have been bought with their contributions. There is an explicit entitlement formula. A retired person's CPP/QPP benefits are determined by a formula, based on age, marital status, and length of working life. The US Social Security system is also similar to a defined benefit programme in this regard.

But there is a crucial difference between private, defined benefit programmes on the one hand, and public plans such as CPP/QPP and the US Social Security plan on the other.

Private defined benefit plans are supposed to (and, with rare exceptions, do) have sufficient funds to meet their obligations. The plans' obligations are the future pension benefits promised to current workers. How do the funds meet these obligations? The easiest way is if the expected present value of future pensions due to a worker exactly equal the expected present value of the payments made to the pension fund on the worker's behalf (by the worker herself, and by the employer). Then all payments could be invested in financial assets. The return on these assets would exactly cover the expected pension benefits owed to the worker when she retires.

Of course, there are complications. The return on the assets is uncertain. So is the present value of the pension payments, since these (typically) will depend on how long the retired worker lives. This uncertainty implies that the pension fund should have assets which exceed, in expected

present value, the pension obligations, to ensure that the fund can meet its obligations even if the market does badly.

Also, a defined benefit plan need not pay each worker a pension which equals, in expected present value, the value of payments made be her (and for her by the employer) while she was working. In a defined benefit plan, there can be, and usually is, some **cross–subsidization**. Some workers subsidize others. But, the plan will be actuarially sound as long as the value of assets held by the pension fund exceeds the expected present value of its aggregate pension obligations.

An alternative to a **funded** pension plan would be an **unfunded**, or **pay as you go** plan. Such a plan works on the following principle. In each year, there are current employees who are paying into the fund, as well as retired employees, who are owed pension benefits. Why not use the **current** payments into the fund, by current workers, to pay the current pension obligations to retired workers?

If the plan is completely unfunded, no assets are acquired by the fund. The current contributions are paid out in benefits to retired people, so that there is no money to invest.

But if current workers' contributions are going to pay benefits to earlier generations of workers (who are currently retired), how will the current workers collect their pension benefits when they retire? There are no assets being acquired by the pension plan, since any money acquired is going to pension payments to current retirees. Future pension benefits, owed to current workers, can be financed from payments made by the next generation of workers. As long as new generations of workers keep coming along, and keep contributing to the plan during their working lives, the plan can keep going along, paying pension benefits to workers of generation t out of payments made into the fund by workers of generation $t+1$.

To a large degree the American Social Security system is a pay–as–you–go plan. The American social security system does hold a large amount of financial assets. However, the value of those assets is considerably less than the expected present value of pension payments which have been promised to current (and past) workers. Canada's plans (CPP and QPP) have a larger fraction of their future obligations backed by assets which they hold. But they're still not fully funded : the assets which are held currently by the CPP administration are not quite large enough to fund the expected future CPP obligations of current Canadian workers.

An unfunded (or incompletely funded) pension plan is bankrupt. So operating a private pension plan (or any other private investment vehicle) on a pay–as–you–go basis is illegal. Unfunded plans are often described as "Ponzi schemes", after a financier from the 1930s. Mr Ponzi promised investors a very high return. However, instead of investing the money he received (as he had promised), he simply paid out money received from one month's new investors to the people who had invested the previous month. As long as the amount of money received each month kept growing, investors were receiving a very good return. But, of course, with a scheme of this nature, eventually the rate of new investment must slow. And once the flow of money coming in slows, the whole scheme must collapse (as it did in Mr. Ponzi's case).

So private firms, and private pension funds, are not allowed to run pay–as–you–go schemes.

These schemes cannot go on forever, since the payments going in must slow down eventually.

The government has a big advantage over private firms here. The money flowing into CPP/QPP, or the American Social Security system, is not being contributed voluntarily. It is coming from taxes. Since taxes seem likely to persist forever, the government could be able to sustain this sort of pay–as–you–go scheme.

What rate of return could a sustainable unfunded public pension plan offer? To simplify matters, consider a world in which each adult lives for two periods : a working period, and then a period (of the same length) of retirement. Of course this is a gross simplification, but it does enable some simple arithmetic results, which can be generalized to a more realistic situation.

Imagine, in this simple world, in which each worker of each generation works for one period and then lives on in retirement for another period, that each worker pays x into a public pension plan during her working life. How large a pension will she be able to collect in retirement? If each generation of workers makes the same contribution, x per person, and if each generation has the same number of workers, then the pension she will collect is the same x that she paid in : it will be the contribution from some worker from the next generation.

So, with a constant contribution rate, and a constant level of population, pension contributions would earn a zero rate of return if the plan was sustainable. Now that may not seem like a very good rate of return. But even a pension plan in which future benefits exactly equal current contributions may have some use. This plan enables workers to reallocate income over their lifetimes. That may be very valuable. People do want to smooth their consumption over their lifetimes.

In this simple model, preferences of people could be represented by some utility function $U(C_Y, C_O)$, defined over consumption when young (and working) C_Y, and consumption when old (and retired) C_O. Typically, we would expect that consumption in the two periods are not perfect substitutes. The curvature of the indifference curve in figure 5 illustrates that people would like to reallocate consumption into a pattern in which they have some consumption in each period, rather than a lot of consumption in one period, and none in the other.

In figure 5, the person's "endowment point" E has also been drawn in. That point is where she would be if she simply consumed what she earned, when she earned it. It is on the horizontal axis, since she earns income only when young (and working). A pension plan which deducts x from her income when she is working, and gives her x when she is retired, moves her up and to the left in the diagram, up a line with slope -1. In the figure, the presence of a pension, even with a zero net rate of return, has moved her to a higher indifference curve, by reallocating income from the period in which she has a lot of it, to a period in which she has none of it.

So, if it were not possible for people to save on their own, an unfunded public pension scheme will make people better off, by enabling people to reallocate consumption over their lifetimes. Suppose, for instance, that there were no productive physical investment opportunities, and that output could not be stored easily. In such an imaginary world, people would have great difficulty reallocating their income from their working years to their retirement years, if there were no public institution. By assumption, they cannot save on their own (because the output has been assumed

not to be easily stored). They would also be unable to make any sort of deals on private markets . What they would want to do is to give up some of their income in their working years, in exchange for some income when retired. But there is no–one with whom such a deal can be made. Today's retired people could certainly use some of this person's income now. But they will not be producing anything to give back when she retires. (They may not even be alive then.) Future workers will have income which they would be willing to trade with this person when she is retired — except that currently they are not even in the workforce, perhaps not even born. In this simple **overlapping generations** model, no bilateral trades between generations are possible. Some sort of long–lived institution, such as a government pension plan, is needed to redistribute among generations.

Now, what was most unrealistic in the previous paragraph's story was the assumption that there was no way for people to save. People can invest. The investment not only enables them to reallocate income from their working years to their retirement years, but it also can earn them a positive return.

However, the advantage of an unfunded public pension plan has also been understated. It was assumed above that the number of workers per generation was constant, and that each worker contributed the same amount each period.

What if the population were growing, at the steady rate n per period? Continue to assume that each worker in each generation pays the same amount x into the pension plan during her working years. How many workers are there in period t? If the population grows at the constant rate n, then there will be

$$((1+n)^t P_0$$

people, if P_0 was the number of workers in period 0. So in period $t+1$, there will be $(1+n)^{t+1}P_0$ workers, and $(1+n)^t P_0$ retired people. (The retired people in period $t+1$ are just the workers from period t.) That means that each retired person can claim a pension, not just of x, but of $(1+n)x$: total collections from workers in period $t+1$ are $(1+n)^{t+1}P_0 x$, so that if these collections are divided equally among the $(1+n)^t P_0$ retired people, then each retired person gets $(1+n)x$.

If the population were growing at the constant rate of n per period, then the maximum rate of return which the unfunded pension plan could offer would be n. But it can offer this return, to each and every worker, without running short of money — as long as the working population continues to grow at the same constant rate.

If productivity is increasing, an even higher return can be sustained. Suppose that each worker in each generation contributes a constant **fraction** of her income into the pension plan. If output per worker were growing at the constant rate a, then each generation's contributions per worker would be $(1+a)$ times the previous generation's workers' contributions. In this case, an even higher rate of return could be sustained. Each retired worker could claim a pension of $(1+n)(1+a)$ times her contributions : that is a rate of return of $a+n+an$. Again, this return can be sustained by an unfunded pension plan only if productivity continues to grow forever at the constant rate a per period.

So an unfunded pension plan can be sustained, and can offer a positive rate of return to pension holders. But the rate is constrained by the growth of the work force, of population, and of productivity. The main reason why Americans perceive their Social Security system to be in "crisis" is not that an unfunded system cannot be sustained : it's that it cannot offer the return that the American social security system has been offering. That is, the formula which determines the benefits to retired people is too generous, given the rate of growth in the U.S. economy. Making the benefit formula less generous would ensure the fiscal viability of the American Social Security system. But it would also mean reneging on a commitment to today's and tomorrow's retired people : the benefit formulae were laws passed by the American Congress.

[Why were the benefits formulae "too generous"? In part, it was a succession of political decisions : promising attractive retirement benefits may get votes, until the voters realize the promises can't be kept. In part, it was a quirk of the procedure used to address inflation in the 1970's and 80's : they "over–indexed" the benefits. In part, it is because the rate of productivity growth in the American economy has slowed, as has the rate of population growth. In large part, it is because the rate of growth of the number of workers contributing to Social Security was very very high up until the 1970's, a rate which could not continue. From its inception in the 1930's, through the 1960's, American Social Security kept expanding its coverage in the U.S. economy. At first only a fraction of industries were subject to Social Security. But this fraction was expanded, until virtually all American workers are covered. Meanwhile, the participation rate of women in the labour force began to grow, increasing the fraction of the population in each generation which was working. Neither of these rates of growth (the proportion of the economy covered by Social Security, and the proportion of women in the labour force) could stay high forever. So Social Security could offer a high rate of return temporarily, while these proportions grew, but could not keep offering these rates forever.]

How does an unfunded public pension plan affect private saving? For simplicity, assume that a person could earn the same rate of return r, on her private saving as on her public pension payments. With no public pension system, what would a person do? She wants to maximize her lifetime utility $U(C_Y, C_O)$, which depends on consumption C_Y when young, and consumption C_O when old. She earns some income Y_Y when young, and perhaps some income Y_O when old. If she chooses to save S, then her lifetime consumption plan will be

$$C_Y = Y_Y - S \tag{1}$$

$$C_O = Y_O + (1 + r)S \tag{2}$$

Each additional dollar she saves will decrease her current consumption by 1, and raise her future consumption by $1+r$. Equations (1) and (2) can be used to define the (C_Y, C_O) combinations which she can afford : she can afford any (C_Y, C_O) such that

$$C_O = [Y_Y(1 + r) + Y_O] - (1 + r)C_Y \tag{3}$$

Equation (3) defines a line, with slope $-(1+r)$ (when C_Y is graphed along the horizontal, C_O on the vertical), through the person's **endowment point** (Y_Y, Y_O). In figure 6, the dotted "opportunity set" is this line. In this figure, I have assumed that the person earns no income when retired, so that $Y_O = 0$, which is why the endowment point is the red dot on the horizontal axis. In the figure, a rate of return of 100% has been assumed ($r = 1$), since the "periods" are supposed to comprise many years.

Able to save (or borrow) at the rate r, the person chooses a consumption point along this opportunity set, by saving (if she wants to move up and to the left along the line). She will choose a consumption plan which gets her the highest level of utility, subject to being on the opportunity set. This is a point at which her indifference curve (between present and future consumption) is tangent to the opportunity set, that is where her MRS equals $1 + r$, which is the slope of the line. In figure 6, the blue dot, labelled "preferred consumption point" is her best choice. In the figure, she achieves that by saving. Her endowment point is $(144, 0)$, and her preferred consumption point is $(96, 96)$. She saves $144 - 96 = 48$, which gives her future consumption of 96 with a rate of return of 100 percent.

Now what would happen if a compulsory public pension plan were introduced? The person's endowment point will shift. First of all, her after–tax income in her working period falls, since she must pay into the public pension. But her income in the retirement period increases, since she now collects some pension income.

As mentioned above, it will be assumed here that the public pension offers the same rate of return as her private savings, r. So if x is her required contribution to the plan, while working, then her working period income falls by x, and her income in retirement increases, by $(1 + r)x$. The compulsory pension plan moves her endowment point : from the red dot at $(144, 0)$ in figure 6, to the green dot at $(108, 72)$. But : as long as the public pension offers the same rate of return as private saving, then the new endowment point is on her opportunity set. The present value of her lifetime income has not been changed.

Since her opportunity set is unchanged, neither is her most preferred point on the opportunity set. She still will want the consumption point at which her MRS equals $1 + r$. She will have exactly the same consumption profile as she would have chosen had there been no public pension : $(96, 96)$ in figure 6. All that has happened is that she has reduced her private saving, since the government is doing some saving for her through the public pension plan. In figure 6, instead of saving 48 to get to the preferred point $(96, 96)$, she only has to save 12 — the government is forcing her to save 36 through the public pension plan.

In other words, a compulsory public pension plan should completely **crowd out** private saving, if the plan offers the same return as does private saving. [If the pension payments were more than the person wanted to save, then she would have to borrow against future pension earnings to get to her preferred consumption plan. That is, if the green dot in figure 6 were to the left of her preferred consumption bundle (the blue dot), then she would want to move right along the opportunity set, instead of left. But if she could borrow at the same rate that she gets when she saves, then the

overall effect is the same.]

If the public pension plan were funded, then it would have no aggregate effect. Every dollar increase in required payments to the pension plan would reduce private savings by 1 dollar. But this reduction in private saving would be offset exactly by an increase in public saving. If the plan were fully funded, the required payments would be invested by the government.

However, if the public pension plan is only partially funded, then there is less public saving to offset the decline in private saving. Assuming still that the public pension offers a rate of return r, the individual will respond in exactly the same manner to an unfunded compulsory pension plan as to a funded compulsory pension plan. Each decreases working period income by x, and increases retirement period income by $(1+r)x$.

That's the main behavioral effect of an unfunded (or partially funded) public pension plan, if individuals are rational, and if the plan offers the same return as private saving : a reduction in overall saving, as private saving is crowded out, with no corresponding increase in public saving (or with less of a corresponding increase in public saving).

If the public pension offered a different return, then the result would be slightly different : the compulsory public pension would move her to a higher (lower) opportunity set if it offered a higher (lower) return. But the main effect would still be a reduction, of more or less the same magnitude, in private saving.

So the main disadvantage of having an unfunded (or under–funded) compulsory public pension plan is the reduction in overall saving. And that is a disadvantage only if there is not enough saving in the economy.

Could there be **too much saving**? Yes, in theory there could be. To see this, let k represent the capital stock per worker in the economy. Let the output produced per worker in the economy be $f(k)$. Since more capital leads to more output, then $f'(k) > 0$. The return to private saving is $r = f'(k)$, the marginal product of capital. But this marginal product declines with the total amount of saving : $f''(k) < 0$.

Now what would be the "best" level of capital stock for a growing economy? I assume that we are trying to determine a long–run plan, where the level k of capital per worker is maintained over time.

If the population is growing, then we will have to keep investing in new capital, just to maintain the same capital–labour ratio k as before. If the economy has a capital–labour ratio of k, and if population is growing at the rate n, then each worker will have to prodcue nk units of new capital, to ensure that next period's workers have the same capital stock — in per capita terms. In this long–run **steady state**, the amount of output per worker is $f(k)$, but nk units of that output have to be set aside to provide capital for the added workers arriving in the next generation. The net output available for current consumption, per worker, is the total output per worker, minus the output which has to be used for investment : $f(k) - nk$.

That means that the best level of capital per worker k, is the level which maximizes $f(k) - nk$. Maximizing $f(k) - nk$ with respect to k implies an optimal capital level (per worker) of k^*, such

that

$$r^* = f'(k^*) = n \tag{4}$$

Equation (4) is referred to as the **golden rule of economic growth**. [It applies only if the economy's technology will lead it to a steady state, in which all variables are constant in per capita terms. If there is enough technological progress, and/or increasing returns to scale, then the economy can keep on growing in per capita terms. The "new" growth theory discussed in (most) macroeconomics courses concentrates on this sort of phenomenon. The golden rule derived here is very much a product of the "old", steady–state theory of economic growth.]

The golden rule of economic growth was derived from choosing the best possible level of capital, without taking into account the cost of acquiring the capital. The only cost considered was the cost of maintaining the capital. A more realistic question would be to find the best level of capital, taking into account that increasing our capital stock does require more saving (and thus less current consumption). That is a more difficult question to solve, since we must find the optimal amount of saving over time, as we increase our capital stock. But — if we do not discount the well–being of future generations compared to the current generation — the solution to that more realistic question is that we should keep saving, gradually accumulating capital, until our capital stock per worker approaches k^*. So the golden rule is a fairly general one. [Although it must be modified if we take into account complications, such as population growth, the fact that capital depreciates over time.]

So, if $r > n$, then our capital stock is less than the golden rule level. That is, it is too low. So, as long as the return to investment exceeds the population growth rate, then savings should be encouraged. Since an unfunded public pension plan discourages saving, it moves the economy in the wrong direction, away from the golden rule level of capital, rather than towards it.

So, should we convert the CPP/QPP to a fully funded system? If our capital stock per person is smaller than what the golden rule suggests, then such a move (towards a fully funded system) would be good for the long–run steady state of the economy. However, we have to get to the steady state. And it certainly is true that moving towards a fully–funded system must cause some difficulty in the short run. If we are going to make the Canadian system fully–funded, then we have either to increase the plan's assets, or reduce its liabilities. Increasing assets means making some group of workers pay more into the system. Decreasing liabilities means reducing pension benefits to some group of retirees. Either way, there will be some people made worse off by the transition.

These transitional costs are inevitable, in going from an unfunded to a fully–funded system. Why did the Americans start with an unfunded system? Starting an unfunded system enables the government to pay pensions immediately to current retirees, even though these retirees did not pay anything into the system. This immediate availability of money for current retirees would not have been possible if they had started up with a funded system,. Then, the first payments into the system would have been saved, rather than paid out. The Americans started up their social security system during the depression of the 1930's, so it was very important to have money

for immediate payout : many older people had been laid off, never to find work again, and many of them had virtually no savings. So short–run distributional considerations made an unfunded system attractive. It offered an immediate "bonus". But any later transition to a funded system must pay for that bonus somehow. That's a question the Americans may have to answer, and which they seem unwilling to answer. The most important question for them (and perhaps for Canadians sometime int he future) is probably not what system to adopt in the long–run. The question probably is "how should the unavoidable cost of that transition to a funded system be shared among the different generations?"