Paper presented at the John Deutsch Institute Conference, "Retirement Policy Issues in Canada", October 26-27, 2007 at Queen's University. The final version is published in a book entitled *Retirement Policy Issues in Canada*, edited by Michael G. Abbott, Charles M. Beach, Robin W. Boadway and James G. MacKinnon, 2009 (Kingston: John Deutsch Institute, Queen's University). Published in cooperation with McGill-Queen's University Press and available at: http://mqup.mcgill.ca .

Some Macroeconomic Effects of Population Aging on Productivity Growth and Living Standards

William Scarth*
McMaster University
October, 2007

Abstract

The prospect of population aging as the baby-boom generation retires has generated much debate. The assessment of this development ranges from a "crisis" to "much ado about nothing," and some analysts have gone so far as to argue that we can expect material living standards to rise, not fall. This paper surveys several approaches to analyzing the growth process to see whether the available literature allows us to narrow down the range of dispute concerning population aging. We focus on relatively simple versions of each approach – ones that can be readily calibrated and applied for direct policy application. We conclude that the aging population does represent a challenge, and those who are concerned about it are not making "much ado about nothing." There is some "good news," however, in the sense that the magnitude of the hit to living standards may just be manageable. It is discouraging, however, that not all the approaches to growth provide analytical support for all policy initiatives that are being discussed as responses to population aging.

*Prepared for the John Deutsch Institute Conference on Retirement Policy Issues in Canada (October 67-27, 2007). Financial support from the SEDAP Research Program and the Dean of Social Science at McMaster is gratefully acknowledged. Comments on this first draft are welcome: scarth@mcmaster.ca

Some Macroeconomic Effects of Population Aging on Productivity Growth and Living Standards

William Scarth
McMaster University
October, 2007

1. Introduction

This paper focuses on a basic worry that underpins many of the specific issues that have been addressed by other papers at this conference. This basic concern is that – with fewer workers, and more individuals in retirement, as the baby-boom generation ages – there will be too few people producing the goods and services that the entire population will want to be consuming. This development is sometimes called a "crisis" (World Bank (1994)), a "crucial" issue (Mintz (2004) or a "demographic storm that we should worry about" (Fortin (2006)). One of the first in Canada to flag this concern was our Auditor General; the 1998 Report concluded that

> "Unless our productivity somehow increases significantly or patterns of work and retirement change substantially, the current demographic trends suggest that the growth in the economy ... will tail off in the coming decades."

No one seems to be prepared to argue that baby-boomers will choose to postpone retirement in a dramatic fashion or that immigration rates will rise dramatically. As a result, the following question emerges as central: Is there anything about the aging population phenomenon itself that can be expected to lead to higher productivity growth? If so, we might be able to defend a more optimistic reaction to the demographic trend. Such optimism would need to be tempered, of course, given our limited inability to explain the productivity slowdown that began in the 1970s. And, while economists' understanding of the growth process has increased a great deal in recent decades, we are still some distance from being able to be confident about offering specific advice to policy-makers. It is interesting that some economists have expressed solid optimism on this broad question, suggesting that population aging may raise living standards. For example, Emery and Rongve (1999) refer to the pessimistic view as "much ado about nothing," and in an article with a title that includes “a positive view on the economics of aging,” Merette (2002) argues that

> "the negative impact of aging on growth … need not be severe … we could even surprise ourselves with … strong performance as our society becomes older."

It must be frustrating for non-specialists to confront an economics profession with views that encompass such a wide spectrum of views regarding the implications of the aging population for living standards. The purpose of this paper is to provide a brief overview of how economists have addressed the question of aging and growth. The hope is that readers may then be able to reach an independent and informed decision concerning the likely impact of aging on material living standards. Should we embrace the dominant view – one of pessimism? Or should we support the minority view – that cautious optimism may be warranted?

The remainder of the paper is organized as follows. In the next section, the traditional approach to understanding economic growth is reviewed. In the following three sections, alternative versions of “new” growth theory are outlined. In each case, the implications demographic developments are explained and, in many cases, the results from numerically calibrated versions of the theory are reported. Concluding remarks are offered in section 6.

2. Traditional Growth Theory

All modern analyses of the growth in living standards start from the base provided by Solow (1956) and Swan (1956). Their simple framework stresses that material well being is limited by what the nation produces, so the initial focus is on the economy's overall input-output relationship – what economists call the production function. It is assumed that output depends on three things: the quantity of labour services employed, the quantity of physical capital (such as machines) that each employee has to work with, and the level of technical knowledge. The model takes the rate of growth of both technical knowledge and the population as given (determined by considerations that are not explained within the model). In addition, it is usually assumed that everyone is employed, so there is no difference between total employment and the population. As a consequence, there are only two simple propositions that drive the model's properties. First, it is assumed that households save a given fraction of their income (the nation's total output). Second, it is assumed that firms incur no costs in taking output that is not purchased by households for current consumption purposes, and converting this output into machines (that are then used – with labour – to produce more goods in future years).

The model’s properties can be appreciated by considering the two standard thought experiments: First, what happens to living standards (per-capita consumption) if households choose a higher savings rate? This question is relevant, given the focus of this paper, since one of the causes of the aging population is rising life expectancy, and an anticipated longer retirement can lead individuals to save more when they are younger. Second, what happens to living standards when there is a lower population growth rate? This question is relevant since an old population has a smaller number of individuals in the child-bearing years. In this paper, I consider a third issue: What happens to living standards when the proportion of the population that is not working rises? To answer these three questions, economists focus is on what is called a balanced growth path – a situation in which all aggregates (total *GDP*, consumption, the capital stock, the quantity of effective workers) grow indefinitely at the same rate. This is an equilibrium, since along such a path, various ratios – such as the capital/labour ratio – stay constant.

Whether more saving is good or bad for living standards depends on how much society is already saving. It can be the case that saving is already so high that we have a particularly large capital/labour ratio. In this case, a lot of each year's output must be put aside to ensure that the capital stock grows as rapidly as does the population. As a result, there is not much currently produced output left over for current consumption, and increased saving only makes this problem worse. But it can be shown that – as long as total profits exceeds total investment spending (a situation that is true for *all* countries in *all* years of observation, see Abel, et. al. (1989) and Scarth (2007b, p. 220)) – we can be confident that our country has not become overly capitalized in this sense. In the empirically relevant version of the model, then, more saving is "good" – in a short-term pain but long-term gain sense. Initially, a higher savings rate must come at the expense of lower consumption. But, as the ensuing increase in the nation's capital stock is brought on stream, labour has more capital to work with, total output per person rises, and a higher *level* of per capita consumption is eventually enjoyed. However, since the ongoing rate of increase in technical knowledge is assumed to be independent of anything that goes on within the model, the *ongoing* rate of increase in living standards is *not* affected at all. Both the short and long run are illustrated in Figure 1.

Figure 1 Effects on Per-Capita Consumption (*C*) of Higher Savings

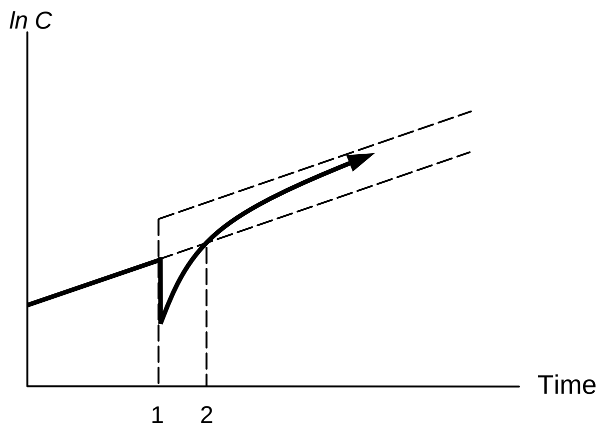

The once-for-all increase in the savings rate takes place at point 1 along the time axis. Had there been no change, per capita consumption would have continued along its exponential growth path (shown here – on a logarithmic scale – as the lower dashed upward-sloping straight line). Initially (between points 1 and 2 in time) the short-term pain emerges, as higher capital accumulation comes at the expense of current consumption. But eventually (beyond point 2 in time), there is enough extra capital to so improve labour's productivity, that overall output, and therefore consumption, is higher. As already emphasized, the *ongoing* growth *rate* of living standards – the slope of the per-capita consumption time path – is *not* affected permanently. So in this strict sense, a pro-savings event (such as increased longevity stimulating a desire for a larger retirement fund) does not permanently raise the growth rate. But if all we want to focus on is a period of 30 years or so, it may not be worth stressing this point too much. This is because the growth path for living standards in Figure 1 *is* steeper during the transition between the two balanced-growth full equilibria.

How does an increase in society's dependency ratio affect the per-capita consumption time path? Answer: it shifts it down in a parallel fashion. Demographers tell us that we can expect about a 10% increase in the overall dependency ratio over the next 30 years. This simple model leaves out considerations such as the particularly high medical expenses that are associated with an older population. Thus, the prediction is that living standards can be expected to fall – in a once-for-all fashion – by about 10%. If readers would like to consider this one-time-level outcome in terms of a growth-rate-equivalent effect, we can note the following. The ongoing annual growth rate of living standards would have to fall by one-third of a percentage point (from an initial annual value of 2%) for 30 years, for living standards to be 10% lower at that 30-year mark. So the base-line growth model predicts a serious drop in living standards as a result of the rising dependency ratio.

This prediction can be considered an under-estimate of the overall effects of aging, since (as already noted) the increased importance of higher medical expenses that will necessary for an older population have been left out. Similarly, the evidence that worker productivity falls as employees age ((Guillemette (2003)) has also been ignored. On the other hand, the numerical prediction given in the last paragraph can be considered an over-estimate, since the possible increase in the savings rate that could accompany higher longevity has been excluded. To include this consideration, we need to assume that households are more purposeful. Instead of simply positing that they follow a fixed rule of thumb (always saving a fixed proportion of income), we need to work out what constitutes optimal planning if households can see the demographic change coming in advance. The literature has addressed this issue. It turns out that – *if* future generations are valued on a par with those currently alive (so that – despite the overlapping generations of an ever larger number of households – decisions are made from the point of view that family dynasties live forever) then the core result that was just discussed goes through with *no* modification. Living standards fall by 10% as the dependency ratio rises by this amount. On the other hand, if future generations and immigrants are not so "loved," individuals can be expected to react to the prospect of their working for a smaller proportion of their lives by choosing to save more in their young years. This higher saving leads to a balanced growth path in which workers have more capital to work with, and this partly mitigates the effect of the rising old-age-dependency ratio. Scarth and Jackson (1998) have shown that the net effect on the level of living standards is a more modest 4% reduction in this case. But as Mankiw (2005) has recently remarked, many individuals (even economists) answer interviews to the effect that the most important thing in their lives is their children. Given this, it is not clear that we should put too much emphasis on this less-loved-future-generations version of the model. In short, it is likely that the increased savings effect is small.

There are other ways to alter the model that have the effect of reducing the magnitude of the predicted hit on living standards that accompanies population aging. One is to note that baby boomers will eventually vanish entirely, so the dependency ratio should eventually fall again. Scarth and Souare (2002) have pursued this issue, and it results in the estimated loss in living standards being pulled down a little more – equivalent to about a one-time loss of about 3%. As already noted, another consideration is that the population growth rate will be lower when the population is older, since old people don't have children. In the Solow-Swan analysis, this is "good news" for living standards. The reasoning is straightforward. The capital/labour ratio needs to remain constant for the economy to be in a balanced growth equilibrium. If the denominator of that ratio is increasing at a slower rate, equilibrium requires a smaller growth rate in the capital stock (the numerator of the capital/labour ratio). Thus, lower population growth frees up a larger proportion of each year's newly produced output to be used for current consumption instead of accumulating capital. This freeing up permits higher living standards, and this fortuitous effect of aging acts as a partial counter-balance for the depressing effect that stems from the higher dependency ratio.

The study undertaken by the Auditor General’s office argues against our attaching too much significance to this favourable effect of a lower population growth rate. Its authors draw attention to the fact that the part of the population that earns lower incomes depends heavily on government transfer payments. With lower population growth, *GDP* growth is lower. Since *GDP* represents the government’s tax base, other things equal, governments will have less revenue with which to finance these transfer payments. Thus, those living hand-to-mouth are hurt by lower population growth. Scarth (2007a) has considered a model which includes both forward-looking households who save for the future (who are helped by a lower population growth rate) and hand-to-mouth households (who are hurt by lower population growth). It is shown that lower population growth is desirable according to the hypothetical compensation criterion, since the “rich” are helped by more than the “poor” are hurt. But this reassurance should be interpreted as having only limited appeal. After all, with the increased international mobility that is available to the owners of capital with globalization, it is increasingly challenging for governments to use the tax system to redistribute income toward the hand-to-mouth group within society. All taxes are ultimately paid by unskilled labour if that group cannot migrate to lower-tax jurisdictions while skilled workers and capital can. Scarth (2007b, p. 195-201) has argued that there is unwarranted pessimism among anti-globalization protesters concerning the scope for governments in small open economies to provide low-income support measures, but none the less, it is a mistake to take this challenge lightly. In any event, many policy-makers find the hypothetical compensation principle to be of limited appeal if fiscal policy is not likely to be used to turn hypothetical into actual compensation. The final reason to temper our enthusiasm concerning how much lower population growth may compensate for the higher-dependency-ratio aspect of aging is that lower population growth has quite the opposite effect on living standards in one of the “new” growth models that we consider in later sections of this paper.

We have just mentioned how adding a “needy” group of households and government transfers to the standard growth model can affect its conclusions. There is another dimension of government that has been stressed in the literature on aging. Merette (2002) and others have argued that Canada’s *RRSP* program is an important consideration. In recent years, this program has provided tax breaks for the baby-boom generation. When this group retires, these individuals will have to accept these accumulated savings as current income, and the amount by which this program deprives the government of revenue will be reduced. Other things equal, this development will make it possible for the government to cut tax *rates*. If the tax on “interest” income is thereby reduced, there will be an increased incentive for households to save. The higher saving can boost living standards, so this is a dimension of the aging population that supports optimism. We discuss the likely empirical significance of this consideration in the next section of the paper, and reach the conclusion that it is quite limited. In the meantime, we summarize the properties of the Solow-Swan approach as follows.

Several aspects of the aging population will lower living standards:

> (i) the higher old-age dependency ratio,
> (ii) the increase in tax rates that will be necessary to maintain the public pension and health care programs (equal to about 5 percentage points of *GDP*),
> (iii) the reduction in labour productivity as individuals age.

Several other dimensions of the aging population will raise living standards:

> (iv) the lower population growth rate,
> (v) the higher savings rate,
> (vi) the reduction in tax rates that can be expected as the *RRSP* system involves a smaller loss in government revenue.

We wish to emphasize two things. First, from an empirical relevance point of view, points (i), (ii) and (iv) are the dominant considerations among the six items in the ledger. Second, point (ii) dominates point (iv) quantitatively. As a result, we can have some confidence that we are not being unduly pessimistic if we assume that that these items cancel off and we focus on point (i) alone. This is because, by doing so, we are under-estimating the severity of the hit that living standards will take as a result of population aging. But, of course, this conclusion is only the verdict derived from traditional growth analysis.

Before moving on to “new” growth theory, it is worth stressing the one key feature of traditional growth analysis that forces that analysis to predict that higher saving can have only a *transitional* effect on the *growth rate* of living standards. This is best appreciated by focusing on households who optimally plan the future time path of their consumption. The standard analysis involves an ever-lasting family dynasty (Ramsey

(1928)) that maximizes utility over time. The simplest function that imposes diminishing marginal utility in any one time period is the present discounted value of the logarithm of consumption for all periods into the future, with a rate of impatience (or time preference) applying at a constant exponential rate. To be maximizing utility, such households must arrange their affairs so that their consumption growth rate equals the excess of the rate of return that can be earned on capital over their rate of time preference. Representing this optimizing condition in simple equation terms, we have:

$$n = r(1-t) - p$$

where *n*, *r*, *t* and *p* represent the growth rate of living standards, the pre-tax interest rate (marginal product of capital), the tax rate, and the rate of impatience. The intuition behind this behavioural rule is straightforward. It pays households to save only if what the market offers them as compensation for forgoing consumption (the after-tax yield on saving) exceeds the discomfort households feel by postponing consumption. So, if the right-hand side of this equation is positive, households do save, and it is the saving that makes positive growth in consumption possible.

Traditional economic analysis specifies that tastes, technology and policy are all determined exogenously (on the basis of outside considerations), and our models are used to determine the remaining items – the resulting economic outcomes. In the present case, parameters *p* and *t* specify tastes and policy (respectively), and – in the traditional growth analysis – the productivity growth rate is assumed to be exogenous as well. As a result, the only variable that can adjust to permit households to behave optimally is the pre-tax return on capital, that is, (variable *r*). To see how things work, consider a pro-savings initiative – a lower tax rate. Initially (before any change in *r*), this policy raises the net return on saving, $r(1 - t)$. So people save more and the nation acquires more capital. But as more and more capital is accumulated, with no corresponding increase in the number of workers, each unit of capital become less productive. (As a simple but telling example, we know that a farmer's rake is not very productive without a worker to use it.) Thus, as time unfolds following the reduction in the tax rate, the pre-tax interest rate gradually falls. The adjustment process is complete when r has fallen by the same amount that $(1 - t)$ has risen. So it is the assumption of diminishing marginal productivity of the man-made input to the production process (capital) that is the driving force behind the proposition that the tax cut cannot permanently raise growth. Modern growth theory is based on the recognition of this fact. All "new" models involve the assumption of constant returns somewhere within the analysis. With this as background, we now consider three versions of new growth theory.

3. "New" Growth Theory: Human Capital

One revised approach for analyzing growth involves the proposition that newly produced output can be used in three ways, not just two. Current production can take the form of consumption, additions to the physical capital stock, or (the new option) additions to the stock of knowledge (human capital). In equation form, this proposition is given as

$$Y = C + \Delta K + \Delta H$$

where the variables are: Y – output, C – consumption, K – physical capital, and H – human capital. The Δ symbol stands for "change in". The production function is

$$Y = K^{a}((1-q)H)^{1-a}$$

where a is a positive fraction and q denotes the proportion of the population that is not working. (If the individual is not at work, her human capital is not available to the production process.) The expression defining the marginal product of each form of capital follows immediately from the total product function. Profit maximization involves firms hiring each factor up to the point that the marginal product equals the rental cost that must be paid to households for using their capital. Hence, the following two equations are part of the model: $aY/K = r$ and $(1-a)Y/(1-q)H = w,$ where r and w are the rental prices for each factor of production.

Optimization at the household level involves two propositions – that the growth in consumption equal the excess of the after-tax yield over the rate of impatience (as explained above), and that the after-tax yield on both forms of capital be the same. These outcomes can be summarized as $\Delta C/C = r(1-t) - p$ and $r(1-t) = w(1-q)(1-t).$ When the equal-yield and the optimal hiring rules are combined, we have $(H/K) = (1-a)/a.$ This outcome can be substituted into the production function, so that it can be re-expressed as $Y = AK$ where $A = ((1-a)/a)^{1-a}(1-q)^{1-a}.$

Given the assumption of balanced growth, the earlier relationships can be summarized as

$$n = \Delta C/C = \Delta K/K = aA(1-t) - p.$$

Since nothing on the right-hand side of this equation changes through time, the prerequisite for an endogenously determined full-equilibrium growth rate is satisfied here (This is because $\Delta K/K$ equals a constant, not an expression that falls in value as K rises.) As a result, the economy's growth rate is permanently affected by any once-for-all changes that are assumed to take place in the variables that appear on the right. Of particular interest to us is the predicted increase in q (from 0.475 to 0.525). For reasonable values of the other parameters, this increase in the dependency ratio permanently reduces the growth rate of living standards by a lot – by just under two-thirds of a percentage point. In terms of a graph like Figure 1, the per-capita consumption time path *pivots* down to a lower slope, and there is no discrete jump (either up or down) at the point in time when this pivot occurs. This outcome is a reduction in living standards that is about twice as big as we obtained when we based the analysis of the increased dependency ratio on the Solow-Swan analysis.

It is important to have some feel for the sensitivity of this predicted outcome to changes in the model's specification. We consider several. First, as above, we have abstracted from the fact that an older population involves higher health-care costs – that have been estimated at about 5 percentage points of *GDP* (Finance Canada (2003)). To this extent, our illustrative numbers are an under-estimate of the true outcome. Second, the particular model outlined here assumes that all future generations are loved by the generations that are currently alive just as much as would be the case if the current generation never died. Hu (1999), Futagami and Nakajima (2001) and Rafique (2006) have considered more general cases in this regard, thereby allowing for an increased savings effect. In addition, they consider longer life expectancy and lower population growth. They ignore the item that is our major focus – the direct effect of a higher dependency ratio. They find that the favourable lower-population-growth and higher-savings effects combine to yield a higher growth rate equal to about six-tenths of a percentage point. It turns out that almost all of this total outcome follows from the drop in the population growth rate, not the small rise in savings.

What about induced changes in tax rates? To address this issue, we focus on the equations that determine n and A – given earlier in this section. The total differential of these relationships yields:

$$\Delta n = -r[((1-t)(1-a)/(1-q))\Delta q + \Delta t].$$

To insulate the growth rate from the rise in the dependency ratio (to ensure $\Delta n = 0$), we require a tax-rate change equal to

$$\Delta t = -((1-t)(1-a)/(1-q))\Delta q.$$

Illustrative initial parameter values are: a = 1/3, t = 1/4 and q = 1/2, and the demographers tell us that a Δq = 0.05 can be expected. To keep the productivity growth rate from falling, then, we need a tax rate cut of 5 percentage points. Merette's (2002) analysis indicates that a tax cut of just this magnitude can be expected (if only the *RRSP* effect is considered). Brown (2002) offers partial support by estimating that a 3 percentage point reduction in tax rates will be possible as a result of the *RRSP* effect. But Finance Canada (2003) has pursued this issue in greater detail, and they argue that a tax rate cut that is much smaller – one-seventh of one percentage point – is all that can be expected. In their sensitivity tests, the biggest tax-rate cut they estimate is 1.25 percentage points. In addition, we must remember that there is the competing effect on the level of tax rates – the need to finance higher public pension and health expenditure when baby boomers are old. As noted, these expenditures will require about 5 percentage points of *GDP*. Thus, there is no way that we can expect tax-rate cuts at all, let alone by the 5 percentage points that we have seen would be necessary to keep the growth rate from falling in the face of the rising dependency ratio. Indeed, the rising tax rates will be enough to cancel out the favourable effects of a higher saving rate and a lower population growth rate. As in the previous section, then, it seems reasonable to argue that we will be erring a little on the too optimistic side if we proceed with the rough-and-ready assumption that all effects other than the *direct* influence of higher dependency cancel out. In short, this endogenous-growth-rate analysis supports the pessimistic view on aging and growth.

There is one further change in the model that could raise the likelihood that this conclusion could be reversed. The economy can be specified with two distinct parts: a manufacturing sector that produces the consumer goods and the physical capital, and an education sector that produces knowledge (the human capital). With distinct sectors it is possible to have two different production functions. The standard approach is to assume a Cobb-Douglas production function in the manufacturing sector: $Y = K^{a}(f(1-q)H)^{1-a}$. This input-output function is the same as we considered above except now there is a proportion *f* which denotes the fraction of the existing human capital that is employed in this sector. The big difference is the assumption that physical capital is not used in the production of knowledge, so the production function in the education sector is $\Delta H = B((1-f)(1-q)H)$. According to this specification, the marginal product of using existing knowledge to produce new knowledge, *B*, does *not* diminish as more knowledge is acquired over time. With the crucial linear growth-rate relationship (needed to have the growth rate permanently affected by developments within the model) embedded in the education sector in this case, it is human capital, not capital in general, that is the "engine of growth." The equal-yields condition is now $r(1-t) = B(1-t)(1-q)$, and the consumption-growth-rate equation becomes $n = B(1-t)(1-q) - p$. For illustration, we calibrate as before: initial $r = B(1-q) = 0.12$ and $\Delta q = .05$. This exercise yields the same estimated effect on the growth rate as we obtained in the one-sector model. However, there may be increased scope for a lower tax rate to counteract this downward pressure on the growth rate in this setting because households have an additional margin of substitution in this case. There is an increased incentive for individuals to invest more in education, and this sector is the engine of growth. There may be less tax breaks offered on wage income (as opposed to income that is derived from physical capital), so there is a better chance for tax-rate increases to be smaller in this two-sector setting. As a result, it is, perhaps, not surprising that the most optimistic estimate concerning aging and growth come from Merette's study that is based on this two-sector structure.

One final issue might concern readers: the models we have surveyed concern a closed economy – one that has no trading relationships with other countries. Canada is an open economy, so the justification for applying these models to Canada lies in the presumption that the aging-population phenomenon is likely to be quite similar in the United States, and it is customary to use a closed-economy framework for analyzing North America as a whole. We have followed this convention. For readers who find this convention unappealing, it is important to note that open-economy analyses have been derived (Lackie (2007)), and the results support the conclusion that aging can be expected to lower productivity growth by a very similar amount.

4. "New" Growth Theory: Research and Development

The empirical applicability of the human capital approach to growth has been challenged. As we discovered in the previous section, calibrated versions of these models predict large growth-rate effects following relatively modest changes in tax or savings rates.

These predictions do not fit well with cross-sectional evidence. For example, compared with other OECD countries, the United States has low tax and savings rates. Indeed, the national saving rate there has been essentially zero for some time. Yet the United States has the highest productivity growth rate. Another undesriable feature of the human capital models is that it is hard to defend the assumption that all knowledge is inextricably tied to workers. Surely it is the case that, after knowledge about a new invention becomes widespread, all participants in the economy have the ability to exploit this knowledge. This is not made impossible simply because some individuals retire. The research and development approach is not open to either of these criticisms.

Aghion and Howitt (1992) have developed a version of the R and D approach that is known as the Schumpeterian paradigm, since it highlights the “creative destruction” that is involved in the innovation process. Their approach formalizes the idea that R and D is carried out by profit-maximizing entrepreneurs, and it involves shifting the location of the linear differential equation (that is necessary for endogenous growth) from the accumulation of capital to the definition of the growth in technology. The rate at which new inventions emerge from the R and D process is proportional to the amount of resources that are devoted to research. An interesting normative analysis of growth policy is possible within this framework since each new invention involves two externality effects – it creates spillover benefits for other producers after the patent period is over and it creates spillover costs in the form of lower profits for the inventor of the previous invention that has then been made obsolete.

It is fortunate that Aghion and Howitt (2007) have recently outlined a more accessible version of the Schumpeterian model. Despite the fact that it involves an extension (the addition of physical capital) to their earlier analysis, it is simpler by virtue of the fact that the inter-temporal-optimization underpinnings for the household consumption-savings choice are dropped and replaced by the Solow-Swan specification – a constant savings rate. This change leads to two appealing features; first, lesser mortals such as the present author can readily follow the derivations, and second, the resulting model is a hybrid of the Solow-Swan and the Schumpeterian approaches. In the hybrid, capital accumulation proceeds as it does in the original neoclassical model, but productivity growth arises endogenously as in the creative-destruction framework. Within this setting, the Solow-Swan model emerges as the limiting case of modern growth theory in which the marginal productivity of efforts to innovate is zero.

Readers are encouraged to consult Aghion and Howitt (2007) directly, since it is not appropriate to repeat their analysis in detail here. Instead, I simply report on one straightforward extension that I have pursued. Their model does not permit analysis of an increase in society's dependency ratio since it is assumed that this item is always unity. Instead, using the notation from earlier sections, I replace the unity value with proportion $(1 - q)$ in all the appropriate equations. After re-expressing the system in change form, I calculate the expression that defines how much the growth rate is affected by an increase in the dependency ratio. Finally, after inserting representation parameter values into this expression I conclude that the annual growth rate can be expected to fall, but only by a very samll amount – by just one-twenty-fifth of one percentage point – as the dependency ratio increases. Aghion and Howitt have stressed the fact that growth is affected only very slightly by changes in the savings rate in their framework. My extension indicates that the same conclusion applies to the effect of aging on the growth rate. The "good news” aspect of this finding is that this branch of modern growth theory provides some support for a “don’t panic” reaction to the prospect of population aging. The “bad news” aspect is that the Schumpetarin approach provides less support for the notion that pro-savings initiatives can be expected to offer significant help in trying to keep the growth rate up as the population ages. This consideration is central to the Canadian policy debate, since the government’s rationale for establishing the 20% debt-ratio target – to be reached by 2020 – is based on the effects of population aging. Two reactions to this concern are possible. First, we can argue that we will not need debt policy to provide a significant cushion for living standards when the baby boomers retire, since (according to this model) the threat to the growth rate of that development is small in the first place. Second, we can remind ourselves that the focus on permanent growth-rate effects in endogenous-growth models tends to distract analysts from the one-time consumption-level effects that are still present within these models. Up to this point, the present discussion is open to this critique. But this is easily rectified. We simply note that – in terms of a negative one-time-level effect on consumption – population aging is still a threat to living standards, and debt reduction can still offer significant relief from this threat in this model. The Scarth-Jackson (1998) analysis was the original motivation for linking the chosen debt policy to the population aging phenomenon. Since that study focused entirely on level, not permanent growth-rate, effects, there is no reason for us to react to the permanent growth-rate effects being tiny in the Schumpeterian appraoch, by being complacent about aging.

5. "New" Growth Theory: Population Growth and Limited Natural Resources

The Schumpeterian growth model has its critics as well. Jones (2003) has raised two points. First, the model involves what is known as scale-effect prediction: the bigger is the *level* of the population, the larger is the productivity *growth rate*. This prediction stems from the simple proposition that the more people there are, the more researchers there are, and so more inventions will emerge from the R and D process. The problem is that data does not seem to support the scale effect; there is no systematic correlation showing that bigger economies have higher growth rates. Jones' second concern is what is known as the knife-edge property of endogenous growth theory. If the engine-of-growth equation is not *precisely* linear, all growth-rate effects become temporary – as in the Solow-Swan framework. According to Jones, it is simply not credible to argue that the processes that generate human capital, or new innovations, are exactly constant-returns-to-scale. This is an especially appealing line of argument if another more easily defended constant-returns relationship can be included in the model. Jones argues that this is possible if we focus on the identities that define the population-growth process.

The amount by which the population grows in any period is exactly equal to the number of births minus the number of deaths. Denoting the birth, death and population growth rates by b, d and g, and the total population as P, we have $\Delta P = (b - d)P = gP$, which is a truly defensible exactly-constant-returns relationship from which to derive endogenous growth. While it is not necessary to bring non-renewable resources into the discussion, it is tempting to do so – given the current debate on the perceived trade-off between our environmental and higher-growth objectives. It is particularly tempting since Groth (2007) has done this, by building on an early contribution by Suzuki (1976). Non-renewable resources involve another simple identity. Denoting the amount of the resource used each period by *R*, and the amount of the remaining stock of the resource as *S*, resource depletion is defined by a simple relationship: $\Delta S = -R.$ To have a straightforward analysis, Suzuki and Groth make the standard assumption, which is similar to the constant-savings-rate specification in the Solow-Swan framework. Using *u* to denote the rate of resource utilization (depletion), we have $R = uS,$ and so there is a second constant-returns relationship in the model: $\Delta S = -uS.$

All that remains to be done is to specify the production and capital-accumulation relationships. With proportion (1 – *q*) of the population working and a Cobb-Douglas specification, we have $Y = K^{\alpha}((1-q)P)^{\beta} R^{\gamma},$ where all exponents are positive. Groth argues on empirical grounds that $(\alpha + \beta)$ may well exceed unity, so we focus on this case. Finally, the capital stock grows whenever current output exceeds total consumption (per capita consumption, *c*, times the population): $\Delta K = Y - cP.$ Assuming a balanced-growth equilibrium, and that we continue to restrict our attention to one-time changes in the old-age dependency ratio, we can show that these relationships imply that the growth in per capita consumption is

$$n = ((\alpha + \beta - 1)g - \gamma u)/(1-\alpha).$$

For the remainder of this section, we discuss several implications of this result. But first we point out that Jones' analysis is slightly different. It is more complicated since he spells out a particular process for the investment in innovation, but it is simpler since he excludes any consideration of non-renewable resources. Our conlusions about demographic developments are exactly what follow in Jones''study.

The first point is that ongoing growth in material living standards is possible – despite the fact that production involves a non-renewable resource. A doomsday conclusion has been avoided since the Cobb-Douglas production function involves the assumption that no input is absolutely essential. But even still, sustained growth is not assured. It requires increasing returns with respect to capital and labour $(\alpha + \beta > 1)$, a "large" population growth rate, and a "small" resource utilization rate. In terms of ongoing growth-rate considerations, an environmentally friendly agenda (a small value for *u*) leads to higher, not lower, growth. In this sense, then, there is not a trade-off between our environmental and our higher-growth objectives.

There is a mixed message concerning demographic developments. The "good news" is that the ongoing growth of living standards is not depressed by an increase in the old-age dependency ratio. The "bad news" is that the other dimension of the aging population – a lower population growth rate – has very different effects in this growth model. Since the growing population is the "engine of growth" in this framework, lower population growth pulls the growth of living standards down. Unless we can be confident that policy-makers should put no weight on this approach to endogenous growth, and I do not believe that we can, then we must now acknowledge that the other approaches to growth may have given us a false impression concerning population growth rates. Except for an undesirable income-distribution consideration – those models supported the view that lower population growth dimension of the aging baby-boomer phenomenon should be welcomed. If this last class of models has relevance, however, this view can no longer be sustained.

6. Conclusions

The purpose of the paper has been to survey growth theory, to select simple versions of each approach that can be readily calibrated and applied for direct policy application, and to discover whether a consistent set of conclusions concerning the effect of population aging on material living standards emerges. Since this objective has forced the paper to be rather taxanomic, it does not seem appropriate to repeat specific conclusions here. Instead, it seems more productive to focus on the overall general message. The aging population *does* represent a challenge, and those who are concerned about it are not making “much ado about nothing.” There is some “good news” but it is not that living standards will rise. The good news is that the likely magnitude of the hit to living standards may just be manageable. Our debt-reduction strategy is proceeding well, and with a concerted effort on the tax reform front, more progress in providing the necessary cushion for living standards may be achieved in time. It is discouraging, however, that not all the approaches to growth provide analytical support for all the policy initiatives that are being discussed as responses to population aging.

References

Abel, A., G. Mankiw, L. Summers and R. Zeckhauser, (1989), “Assessing Dynamic Efficiency: Theory and Evidence,” *Review of Economic Studies* 56, 1-20.

Aghion, P. and P. Howitt, (1992), “A Model of Growth through Creative Destruction,” *Econometrica* 60, 323-351.

Aghion, P. and P. Howitt, (2007), “Capital, Innovation and Growth Accounting,” *Oxford Review of Economic Policy* 23, 79-93.

Brown, R., “Paying for Canada’s Aging Population: How Big is the Problem?” Canadian Institute of Actuaries Member’s Paper.

Emery, H. and I. Rongve, (1999), “Much Ado About Nothing: The Demographic Bulges, the Productivity Puzzle, and CPP Reform,” *Contemporary Economic Policy* 17, 68-78.

Finance Canada, (2003), *Tax Expenditures and Evaluations 2003*. http://www.fin.gc.ca/taxexp/2003/taxexp03_e.html

Fortin, P. (2006), “Back to Basics,” *Policy Options* (April/May, 59) Montreal: Institute for Research on Public Policy.

Futagami, K., and T. Nakajima, (2001), “Population Aging and Economic Growth,” *Journal of Macroeconomics* 23: 31-44.

Guillemette, Y., (2003), *Slowing Down with Age: The Ominous Implications of Workforce Aging for Canadian Living Standards*. Commentary No. 182. Toronto: CD Howe Institute.

Groth, C., (2007), “Growth Essential Non-Renewable Resources and Limits to Growth,” University of Copenhagen Working Paper.

Hu, S-C., (1999), “Economic Growth in the Perpetual Youth Model: Implications of the Annuity Market and Demographics,” *Journal of Macroeconomics* 21, 107-124.

Jones, C., (2003), “Population and Ideas: A Theory of Endogenous Growth,” in *Knowledge, Information, and Expectations in Modern Macroeconomics*, ed. P. Aghion, R. Frydman, J. Stiglitz and M. Woodford. Princeton: Princeton University Press.

Lackie, C., (2007), “A Target Debt-to-GDP Ratio with an Aging Population,” McMaster University, MA in Economic Policy, Policy Project.

Mankiw, G., (2005), "Comment," Brookings Papers on Economic Activity, Issue 1, 316-325.

Merrette, M., (2002), “The Bright Side: A Positive View on the Economics of Aging,” *Choices* 8(1). Montreal: Institute for Research on Public Policy.

Mintz, J., (2004), "Is the Debt War Over? What Have We Learned?" in *Is the Debt War Over?* ed. C. Ragan and W. Watson. Montreal: Institute for Reasearch on Public Policy.

Office of the Auditor General of Canada (1998) “Population Aging and Information for Parliament: Understanding the Choices,” *Report of the Auditor General to the House of Commons – 1998*.

Rafique, S., (2006), “Economic Effects of Declining Population Growth,” McMaster University, MA in Economic Policy, Policy Project.

Ramsey, F., (1928), “A Mathematical Theory of Saving,” *Economic Journal* 38, 543-559.

Scarth, W., (2007a), “Lower Population Growth and Average Living Standards,” (mimeo).

Scarth, W., (2007b), *Macroeconomics: An Introduction to Advanced Methods, Third Edition*, Toronto: Thomson Custom Publishing.

Scarth, W., and H. Jackson (1998), “The Target Debt-to-GDP Ratio: How Big Should It Be? And How Quickly Should We Approach It?” in *Fiscal Targets and Economic Growth*, ed. T. Courchene and T. Wilson. Kingston: John Deutsch Institute for the Study of Economic Policy.

Scarth, W, and M. Souare, (2002), “Baby Boom Aging and Average Living Standards,” SEDAP Research Paper No. 68, McMaster University.

Solow, R., (1956), “A Contribution to the Theory of Economic Growth,” *Quarterly Journal of Economics* 70, 65-94.

Suzuki, H., (1976), “On the Possibility of Steadily Growing Per Capita Consumption in an Economy with a Wasting and Non-Replenishable Resource,” *Review of Economic Studies* 43, 527-535.

Swan, T., (1956), “Economic Growth and Capital Accumulation,” *Economic Record* 32, 145-166.

World Bank (1994), *Averting the Old Age Crisis: Policies to Protect the Old and Promote Growth*. Oxford: Oxford University Press.