

DYNAMIC
MACROECONOMIC
ANALYSIS

Theory and Policy in General Equilibrium

edited by

SUMRU ALTUG, JAGJIT S. CHADHA

and

CHARLES NOLAN



CAMBRIDGE
UNIVERSITY PRESS

published by the press syndicate of the university of cambridge
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

cambridge university press
The Edinburgh Building, Cambridge, CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa

<http://www.cambridge.org>

© Cambridge University Press 2003

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2003

Printed in the United Kingdom at the University Press, Cambridge

Typeface Adobe Garamond 11/12.5 pt. *System* L^AT_EX 2_ε [tb]

A catalogue record for this book is available from the British Library

Library of Congress Cataloging in Publication data

Dynamic macroeconomic analysis : theory and policy in general equilibrium / edited by
Sumru Altug, Jagjit S. Chadha, Charles Nolan.

p. cm

ISBN 0 521 82668 3 – ISBN 0 521 53403 8 (pb.)

1. Macroeconomics. 2. Equilibrium (Economics) I. Altug, Sumru. II. Chadha, Jagjit.
III. Nolan, Charles.

HB172.5.D96 2003 339.5 – dc21 2003043589

ISBN 0 521 82668 3 hardback

ISBN 0 521 53403 8 paperback

The publisher has used its best endeavours to ensure that the URLs for external
websites referred to in this book are correct and active at the time of going to press.
However, the publisher has no responsibility for the websites and can make no guarantee
that a site will remain live or that the content is or will remain appropriate.

Contents

<i>List of contributors</i>	page vii
<i>Foreword by William A. Brock</i>	xiii
<i>Preface</i>	xxvii
1 The application of stochastic dynamic programming methods to household consumption and saving decisions: a critical survey	1
<i>James Pemberton</i>	
2 Investment dynamics	34
<i>Fanny S. Demers, Michael Demers and Sumru Altug</i>	
3 Taxes and welfare in a stochastically growing economy	155
<i>Stephen J. Turnovsky</i>	
4 Recent developments in the macroeconomic stabilisation literature: is price stability a good stabilisation strategy?	212
<i>Matthew B. Canzoneri, Robert E. Cumby and Behzad T. Diba</i>	
5 On the interaction of monetary and fiscal policy	243
<i>Jagjit S. Chadha and Charles Nolan</i>	
6 Dynamic general equilibrium analysis: the open economy dimension	308
<i>Philip R. Lane and Giovanni Ganelli</i>	
7 Credit frictions and ‘Sudden Stops’ in small open economies: an equilibrium business cycle framework for emerging markets crises	335
<i>Cristina Arellano and Enrique G. Mendoza</i>	
8 Asset pricing in macroeconomic models	406
<i>Paul Söderlind</i>	

9	Labour market search and monetary shocks <i>Carl E. Walsh</i>	451
10	On the introduction of endogenous labour income in deterministic and stochastic endogenous growth models <i>Stephen J. Turnovsky</i>	487
11	Growth and business cycles <i>Gabriel Talmain</i>	509
	<i>Author index</i>	569
	<i>Subject index</i>	576

*The application of stochastic dynamic programming
methods to household consumption and saving
decisions: a critical survey*

James Pemberton

1 INTRODUCTION

This chapter discusses work which applies the methods of stochastic dynamic programming (SDP) to the explanation of consumption and saving behaviour. The emphasis is on the intertemporal consumption and saving choices of individual decision-makers, which I will normally label as ‘households’. There are at least two reasons why it is important to try to explain such choices: first, it is intrinsically interesting; and, second, it is useful as a means of understanding, and potentially forecasting, movements in aggregate consumption, and thus contributing to understanding and/or forecasts of aggregate economic fluctuations. The latter motivation needs no further justification, given the priority which policy-makers attach to trying to prevent fluctuations in economic activity. The former motivation – intrinsic interest – is less often stressed by economists, but it is hard to see why: it is surely worthwhile for humankind to improve its understanding of human behaviour, and economists, along with other social scientists, have much to contribute here.

The application of SDP to household consumption behaviour is very recent, with the first published papers appearing only at the end of the 1980s. Young though it is, this research programme has already changed significantly the way in which economists now analyse consumption choice, and has overturned a number of previously widely held views about consumption behaviour. Any research programme which achieves such outcomes so quickly would normally be judged a success, and in many respects this is an appropriate judgement here. The judgement needs to be qualified, however, on at least two counts. First, some of the ideas which the SDP research programme has overturned, although previously widely believed by mainstream economists, were never subscribed to by those working outside the mainstream. Non-mainstream economists might argue that the SDP

programme has simply allowed the mainstream to catch up with their own thinking. Second, there is room for doubt that SDP methods really capture at all well the ways in which humans actually make decisions.

These issues are considered in the rest of this chapter. Section 2 reviews the development of economists' thinking about consumption behaviour since the time of Keynes, and places the SDP programme in this longer term context. Section 3 looks in more detail at some of the most prominent contributions to the SDP research programme. Section 4 considers criticisms of the SDP programme, and looks at other possible approaches to modelling consumption behaviour. Section 5 draws some conclusions.

2 THE DEVELOPMENT OF HOUSEHOLD CONSUMPTION MODELLING

2.1 Keynes, Modigliani and Friedman

Modern interest in consumption and saving behaviour started with Keynes (1936). Keynes' emphasis on aggregate demand as a short-run determinant of the level of economic activity required him to consider the major components of aggregate demand, of which aggregate private consumption is easily the largest in typical market-based economies. In order to model consumption behaviour, Keynes introduced a theoretical concept – the marginal propensity to consume (MPC) – which has remained central to all subsequent work on consumption and saving behaviour. The basis of Keynes' modelling of the MPC was his 'fundamental psychological law. . . that men are disposed, as a rule and on the average, to increase their consumption as their income increases, but not by as much as the increase in their income' (Keynes 1936: 96). In the context of modern perspectives on consumption behaviour, a difficulty with this statement is that it does not define the 'income' concept: e.g. is it current, or permanent, income which matters? And does it make any difference whether or not the increase in income was previously anticipated? On the other hand, it is interesting to note that Keynes' accompanying discussion of the MPC (1936: chapters 8 and 9) anticipated a number of issues which have become central to more recent work on consumption: e.g. he allows for windfalls, changes in time preference, changes in expected future income relative to current income, precautionary motives, changing family needs over the life cycle and inter-generational bequests.

It soon became evident that Keynes' model of an MPC between zero and one was at odds with time series evidence showing a roughly constant ratio

of aggregate consumption to aggregate income despite persistent growth in the latter. This issue was addressed by the life cycle hypothesis (LCH) and the permanent income hypothesis (PIH) associated, respectively, with Modigliani (Modigliani and Brumberg 1954) and with Friedman (1957). Virtually all economists since then have accepted the basic idea of both the LCH and the PIH, namely that households are *forward-looking*: they are concerned about future as well as current consumption, and they take account of expected future income as well as current income.

Friedman's original (1957, 1963) statement of the PIH implied a very flexible framework. For example, he argued that (i) when calculating permanent income, the discount rate used to obtain the present value of future income is a highly subjective concept, not necessarily bearing any relation to market interest rates (e.g. Friedman 1963 argued for an annual discount rate of around 33 per cent, implying a time horizon of around three years); (ii) different discount rates may be applied to different types of income; and (iii) permanent income may be expected to vary over the future, as new information is acquired (Friedman 1957, figure 1.2:24). Such a framework allows uncertainty (about future income in general, or about particular sources of income, or about future preferences) to have a large influence on behaviour. The short time horizon implied by point (i) is one possible way in which the framework permits households to respond to such uncertainty. By contrast Modigliani's LCH, as retrospectively summarised in his Nobel Lecture (Modigliani 1986) is more closely circumscribed: the time horizon is the remaining expected life cycle, and choice is governed by 'the self-evident proposition that the representative consumer will choose to consume at a reasonably stable rate, close to his anticipated average life consumption' (1986: 301). This leads to the simple, but very influential, diagrammatic representation of life cycle behaviour in figure 1.1, henceforth the 'Modigliani diagram'. The household's income, OY , is constant throughout working life (OR) and falls to zero in retirement (RN); its 'self-evident' desire to maintain constant consumption leads to the steady consumption level OC and the pattern of asset accumulation and decumulation OAN . Of course, numerous complications (e.g. variable working income, variable household size, etc.) are ignored, but many of these would not alter the basic idea.

The simplified life cycle behaviour illustrated in figure 1.1 became the received orthodoxy of mainstream economics as the LCH became the standard framework. To illustrate its influence and durability, it has survived unchanged through all eight editions of one of the leading international student textbooks on intermediate macroeconomics (Dornbusch, Fischer and




Figure 1.1 The 'Modigliani' diagram

Startz 2001). The only change in successive editions has been the gradual inclusion of more material discussing ideas from SDP models – which contest the whole basis of figure 1.1 – though this is given less prominence than the basic Modigliani diagram. Thus, many world-wide cohorts of economics students have absorbed the ideas of figure 1.1 as the basis of mainstream economics views about consumption and saving behaviour.

Figure 1.1's basic idea can easily be derived from the following LCH set-up:

$$V(t) = E_t \sum_{i=0}^N (1+d)^{-i} u[c(t+i)] \quad (1.1)$$

$$E_t \sum_{i=0}^N (1+r)^{-i} c(t+i) = F(t) + H(t) \quad (1.2)$$

$$H(t) = E_t \sum_{i=0}^N (1+r)^{-i} w(t+i). \quad (1.3)$$

$V(t)$ is the lifetime objective function at time t , $c(n)$ is consumption at time n , $u(c(n))$ is the one-period utility function at time n , and $d > 0$ is the per period rate of time discounting. N is the number of remaining lifetime periods after the present period. E_t is the expectations operator as at time t . Equation (1.1) thus says that the household maximises expected discounted

lifetime utility. Equations (1.2) and (1.3) specify the lifetime budget constraint, with $F(t)$ and $H(t)$ denoting accumulated non-human wealth and human wealth, respectively at time t , $w(t)$ denoting labour income at time t ,¹ and r being the per period real interest rate. (Following much, though not all, of the literature, I assume that both r and d are constant throughout this chapter.) (1.1)–(1.3) provide a reasonably general statement of the lifetime maximising problem. They generate the Modigliani diagram with the following restrictions: no uncertainty about the future; and interest and time preference rates which are equal to one another ($r = d$). The latter is a convenient simplification which affects details rather than fundamental principles,² but ruling out uncertainty is a critical assumption, whose effect is that an increase of X in current income has exactly the same impact on current consumption as an increase of $(1 + r)^N X$ which will occur N periods in the future, no matter how large is N . This simple, but striking, proposition, together with the Modigliani diagram and the model of (1.1)–(1.3) with uncertainty excluded, became the received mainstream economics view about consumption and saving behaviour in the 1960s and 1970s. Note that, as emphasised above, it is very different from Friedman's original version of the PIH model. Thus, the frequent references in the literature to the 'LCH/PIH framework' are quite misleading in their implication that the two models are virtually identical. Friedman's approach is more sophisticated than Modigliani's, but it is not so easily reduced to a simple framework such as (1.1)–(1.3) and figure 1.1. Thus, many of Friedman's sophisticated complications were lost sight of in the 1960s and 1970s as mainstream economics adopted the simplified LCH model.

2.2 Hall and random walks

Hall (1978) started from the model of (1.1)–(1.3), including the assumptions of constant r and d (not necessarily $r = d$), and focused on the implications of uncertainty about future labour income. The centrepiece of Hall's paper was the then relatively unfamiliar, but now standard, first-order intertemporal optimising equation, now usually termed the Euler equation:

$$E_t u'(c(t+1)) = [(1+d)/(1+r)]u'(c(t)). \quad (1.4)$$

¹ Equation (1.3) allows labour income to continue right to the end of the life cycle, but a retirement phase can be allowed by setting $w = 0$ for later periods.

² If r is more (less) than d , then in figure 1.1 consumption rises (falls) over the life cycle rather than remaining constant.

Equation (1.4) indicates that the marginal utility of consumption evolves as a random walk, with a trend if r and d are not equal, and trendless if $r = d$. An important special case occurs if utility is quadratic, i.e.

$$u(c(t)) = Ac(t) - Bc(t)^2, \quad A > 0, \quad B > 0. \quad (1.5)$$

Then marginal utility is linear in consumption, and (1.4) implies:

$$E_t c(t+1) = a_0 + a_1 c(t) + e(t+1). \quad (1.6)$$

When $r = d$, $a_0 = 0$ and $a_1 = 1$; a_0 is increasing, and a_1 is decreasing, in $(r - d)$. $e(\cdot)$ is a random shock, with $E_t e(t+1) = 0$. Thus, consumption itself follows a random walk. This case has become known as the *certainty equivalent* (CEQ) model, because its implications for intertemporal consumption and saving choices are equivalent to those in the basic Modigliani model, despite the presence of uncertainty. In particular, (1.6) implies exactly the same propensity to consume out of (expected) lifetime resources as is implied by (1.1)–(1.3) with perfect certainty. It also implies, exactly as in the Modigliani model, that an increase in current income of X or an increase in expected income N periods in the future of $(1+r)X^N$ have identical effects on current consumption.

Where (1.6) differs is in the random walk response to income shocks, which by assumption do not occur in the Modigliani framework. The basic idea is most easily understood if $r = d$ is assumed, so that $a_0 = 0$ and $a_1 = 1$ in (1.6). Then, if there were no shocks so that $e(\cdot) = 0$ in every period, consumption would be constant over life, equal to permanent income, exactly as in the Modigliani diagram. Under uncertainty it is still optimal to plan for consumption to be constant and equal to permanent income, but now any income shock causes calculated permanent income to be revised upwards or downwards, and planned consumption is likewise revised so as to be constant at the new permanent income level. Thus, consumption's random walk derives from the random walk followed by permanent income. Notice the similarity to Friedman's earlier emphasis on the fact that permanent income is typically revised over time.

This model triggered a large empirical research programme, aimed at testing the prediction that households revise their consumption by an amount equal to the permanent income value of any unexpected income change. The income process is critical here. If unexpected changes in income are quickly reversed, their permanent income value is tiny and consumption should hardly alter. Conversely, if unexpected changes are long-lasting, then consumption should alter roughly in line with the current income shock. Early tests of Hall's model (e.g. Flavin 1981) concluded that consumption

is excessively sensitive to income (i.e. consumption alters by more than the permanent income value of shocks), but later work (e.g. Campbell and Deaton 1989) suggests the reverse, i.e. that consumption is too smooth in relation to income. The difference reflects different estimated models of the income process: Flavin's work suggested that income levels follow an autoregressive process, whereby a rise of X in current income implies a rise of less than X in permanent income; by contrast, Campbell and Deaton's work suggested that the growth rate – not the level – of income is autoregressive, whereby a rise of X in current income implies a rise of more than X in permanent income.

2.3 Beyond the certainty equivalent model

The CEQ model of (1.5) and (1.6) is a special case of (1.4). Quadratic utility is less plausible as an assumption about preferences than a constant relative risk aversion (CRRA) utility function:

$$u(c(t)) = \begin{cases} \frac{c(t)^{1-\sigma} - 1}{1-\sigma} & (\sigma \neq 1) \\ \ln c(t) & (\sigma = 1). \end{cases} \quad (1.7)$$

The CRRA assumption is now routinely used in most SDP treatments of consumption, and is adopted in most of the rest of this chapter. To see its implications, substitute (1.7) into (1.4) and rearrange:

$$c(t)^{-\sigma} = [(1+r)/(1+d)] E_t c(t+1)^{-\sigma}. \quad (1.8)$$

Unlike in the CEQ case, the Euler equation is now no longer linear in consumption. This fundamentally alters both the economic implications, and the technical treatment, of the model. The modern SDP analysis of consumption focuses on these economic and technical issues.

Looking first at the basic economic implications, note that the third derivative of the utility function is zero with quadratic utility, as in (1.5), but is positive with CRRA utility, as in (1.7). Consider a simplified context in which the life cycle consists of just two periods – the importance of this simplification is considered shortly – and in which second-period income can be either 'high' or 'low' with equal probability. Then second-period consumption takes either the 'high' value $c(2H)$, or the 'low' value $c(2L) < c(2H)$, with equal probability. These two possible second-period outcomes are depicted on the horizontal axis of figure 1.2, and $E_1 c(2)$ denotes mean expected second period consumption, viewed from the vantage




Figure 1.2 Second-period marginal utility of consumption

point of period 1. The vertical axis measures the marginal utility of second-period consumption, and the curve XX depicts the relationship between consumption and marginal utility when utility is CRRA. The marginal utilities associated with $c(2L)$, $c(2H)$ and $E_1c(2)$ are shown on the vertical axis. The expected marginal utility is

$$E_1u'(2) = [u'(c(2H)) + u'(c(2L))]/2, \quad (1.9)$$

and is also shown on the vertical axis. The key point is that $E_1u'(2)$ is more than $u'(E_1c(2))$: expected future marginal utility is more than the marginal utility of expected future consumption. This inequality must always hold if the third derivative of the utility function is positive. By contrast, in the CEQ model with a zero third derivative, the curve XX in figure 1.2 is replaced by a straight line, and the inequality is converted into an equality.

Contrasting the linear and non-linear versions of XX in figure 1.2 illustrates why most economists believe CRRA (the non-linear case) to be more plausible than quadratic utility (the linear case). The non-linear case drawn in figure 1.2 implies that as consumption goes towards zero, so the marginal utility of consumption goes towards infinity – implying, surely plausibly, that a destitute individual places enormous value on a small consumption gain – whereas linearity implies that marginal utility is no higher at zero consumption than at any positive consumption level (so that, seemingly implausibly, someone who is destitute does not value extra

consumption specially highly). This is why most economists place greater trust in the results of models using CRRA utility than in results from CEQ models.

2.4 Precautionary saving

The implications of replacing (1.5) with (1.7) can be seen by looking again at the Euler equation (1.4). For a given pattern of uncertainty about future income, there is a particular time path of present and future consumption which satisfies the Euler equation (1.4) under CEQ preferences. With CRRA preferences, this time path cannot satisfy (1.4) because the non-linearity in figure 1.2 raises expected future marginal utility relative to current marginal utility. To restore equality requires a different consumption time path: present consumption has to be lower (which increases current marginal utility), and planned future consumption has to be higher (so as to lower expected future marginal utility), compared with the optimal CEQ time path. Thus, for a given pattern of future income uncertainty CRRA preferences yield lower current consumption, and more saving, than CEQ preferences. This extra saving compared with the CEQ case results from *any* utility function with a positive third derivative: CRRA is simply a convenient special case. The extra saving is often labelled *precautionary saving*. Kimball (1990) provides a framework for analysis. He distinguishes between risk aversion and what he terms ‘prudence’. The standard measurement of the degree of absolute risk aversion is $-u''(c)/u'(c)$; Kimball proposes an analogous measure of the degree of absolute prudence, $-u'''(c)/u''(c)$. The conceptual distinction is that ‘[the] term “prudence” is meant to suggest the propensity to prepare and forearm oneself in the face of uncertainty, in contrast to “risk aversion”, which is how much one dislikes uncertainty and would turn away from uncertainty if possible’ (Kimball 1990: 54).

The idea of precautionary saving predates the recent literature – see, e.g., Leland (1968) – but until relatively recently, no systematic work had been done. This reflected substantial technical difficulties in extending the two-period example in figure 1.1 to a multi-period context. The Euler equation (1.4) relates marginal utility in the current period and the next period, but it does not provide a self-contained solution unless the second-period is also the last period, since otherwise the second-period solution depends on a further Euler equation relating periods two and three; the period-three solution in turn depends on period four, and so on. SDP is the standard means of solving this sort of problem, but until relatively recently

the computing power needed to solve multi-period SDP problems was not available to most researchers.

3 RECENT RESEARCH ON CONSUMPTION

3.1 SDP solutions of life cycle problems

SDP is applied to multi-period life cycle consumption problems by defining a final period of life, T , in which all remaining resources are consumed. This defines a two-period problem between periods T and $(T - 1)$, the solution to which implicitly defines optimal $c(T - 1)$. The latter is in turn inserted into another two-period problem between $(T - 1)$ and $(T - 2)$, from which optimal $c(T - 2)$ emerges, and so on. The appendix (p. 31) enlarges on the methods used; here the focus is on the underlying economic issues.

The first application of this backward induction procedure to a multi-period life cycle problem was by Zeldes (1989), and his computing techniques have been followed by others, albeit with more complicated problems becoming feasible as computing power expanded rapidly during the 1990s. The problem of computing power is well illustrated by Zeldes' own description of the difficulties of simulating a version of his model involving both permanent and transitory shocks to income. He was unable to solve this model over more than fifteen periods because it required 'creating two matrices with about 625,000 elements each. The optimal consumption (and value function) then had to be determined for each of the 625,000 possible nodes, for each of the fifteen periods' (Zeldes 1989, n. 22: 286). This exhausted available computer memory. Earlier researchers lacked the computing power to attempt even this; subsequent researchers have solved progressively more complex problems. Thus, research progress during the 1990s was triggered by the availability of greater computing power rather than by new theoretical ideas. The use of this power, however, has itself generated a number of new ideas and insights, which are outlined in the rest of this section. At the same time the sheer complexity of the problem constitutes a potential objection to SDP methods; this issue is taken up in section 1.4.

3.2 Basic SDP life cycle results

Recall some basic propositions of the Modigliani/Hall LCH framework: (1) the present value of lifetime consumption cannot exceed that of lifetime income, but other than this there is no connection between the two. In particular, the pattern of lifetime consumption is divorced from that of

lifetime income (cf. figure 1.1); (2) the MPC for a current-period increase in income of X is the same as that for an increase of $(1 + r)^N X$ expected to occur N periods in the future; and (3) a household changes its current consumption by the permanent income equivalent of any unexpected change in current income.

SDP analysis using CRRA utility comprehensively overturns these results. Instead, what emerges is the following:

- (1) The MPC is generally larger, often much larger, out of an increase in current income than out of an increase of equal present value in expected future income.
- (2) This disparity is especially important for households whose current accumulated financial wealth is low relative to normal income. Conversely, households with high levels of financial resources relative to normal income behave more like Modigliani consumers.
- (3) A household's planned rate of growth of future consumption is negatively related to its current asset holdings.
- (4) There is no systematic relationship between the MPC out of a current unexpected income change, and the latter's permanent income equivalent.
- (5) The pattern of lifetime consumption is not divorced from that of lifetime income. Instead, consumption closely tracks income, at least for households below the age of around 45–50.

3.3 The buffer stock model

I focus first on Carroll (1997), who provides an interesting and fully developed application of SDP to consumption and saving issues. I then look more briefly at other influential SDP models.

Carroll assumes that the household maximises an objective function of the form of (1.1) (with the horizon N sometimes allowed to go to infinity) assuming CRRA utility as in (1.4). He specifies the relevant budget constraints and income processes as follows:

$$x(t + 1) = (1 + r)[x(t) - c(t)] + E_t w(t + 1); \quad (1.10)$$

$$E_t w(t + 1) = E_t [p(t + 1)v(t + 1)]; \quad (1.11)$$

$$E_t p(t + 1) = (1 + g)p(t)E_t n(t + 1); \quad (1.12)$$

$$v(t) = \begin{cases} 0 & \text{with probability } q \\ Z & \text{with probability } 1 - q; \end{cases} \quad (1.13)$$

$$\ln Z \approx N(-\sigma_{\ln Z}^2/2, \sigma_{\ln Z}^2); \quad (1.14)$$

$$\ln n \approx N(-\sigma_{\ln n}^2/2, \sigma_{\ln n}^2). \quad (1.15)$$

$x(t)$ is ‘cash on hand’ at time t , defined as:

$$x(t) = F(t) + w(t). \quad (1.16)$$

Labour income $w(t)$ has both a persistent component, $p(t)$, and a transitory component, $v(t)$. The persistent component grows at a trend rate of g per period, and is also subject to shocks, n , specified in (1.15) such that $E_t n(t+1) = 1$. Carroll designates the trend path of the persistent income component as ‘permanent income’, but this is potentially confusing because it is by no means the same as the usually understood definition of ‘permanent income’, which is based on the notion of a constant annuity. I therefore label it as ‘normal income’ in what follows.

Disturbances to the persistent income component lead to persistent changes in income. Labour income in any period is also subject to transitory shocks $v(\cdot)$. These are of two types. With probability q they cause labour income to be zero (e.g. unemployment), and with probability $(1 - q)$ they have less dramatic effects, causing income to vary around its ‘normal’ level (e.g. variations in bonuses, overtime payments, etc.).

Before considering the model further, some limitations of it and of most of the other work considered in this chapter should be mentioned. First, it assumes perfect capital markets, whereas many economists might consider it more reasonable to focus on the implications of market imperfections, especially limitations on the ability to borrow against expected future income. This is briefly discussed later in the chapter. A second limitation is that the model assumes only one type of saving asset, whose returns are themselves not subject to uncertainty. Although this assumption is often made in the literature, it is oversimplified in two respects: there are many potential savings vehicles, and most of these are subject to at least some uncertainty of return. Thus, consumers have to contend not only with labour income uncertainty, but also with choices among different saving assets with different risk–return patterns. In practice, focusing only on labour income uncertainty is a sensible strategy in analysing life cycle behaviour at least for consumers in the first half to two-thirds of their working lives. Few such consumers have significant amounts of discretionary financial wealth, so that for them, labour income uncertainty is overwhelmingly the most important source of uncertainty. Introducing capital as well as labour income uncertainty would substantially complicate the model without adding much insight for consumers in this age range. A third simplification is that consumers are assumed to have no source of insurance against future income uncertainty other than self-insurance. Thus, markets offering insurance contracts against such contingencies are assumed absent, as of course is

the case in reality. The obvious reason for this is the problem of moral hazard. Some of these insurance possibilities are, however, provided in many countries by the state, and later in the chapter I look at the importance of such social safety nets: they have an important impact on SDP models. The chapter does not discuss other possible sources of insurance against adverse income contingencies such as intra-family transfers.

Many of Carroll's results are based on the log linearised Euler equation, which takes the form:

$$\Delta \log c(t+1) \cong \sigma^{-1}(r-d) + 0.5(1+\sigma)E_t s^2(t+1) + e(t+1) \quad (1.17)$$

$s(t+1)$ is implicitly defined by $c(t+1) = [1 + s(t+1)]c(t)$. The implications of (1.17) can most easily be understood in combination with what Carroll labels the 'impatience condition', which determines whether or not the optimal consumption rule derived from the model of (1.10)–(1.15) by SDP-type backward induction converges. Optimal consumption in any period t is defined by normalising both consumption and cash on hand by 'normal' labour income $p(t)$ for that period, and specifying the optimal consumption rule for period t in terms of these normalised variables:

$$[c(t)/p(t)] = f[x(t)/p(t)]. \quad (1.18)$$

Carroll shows that the function $f(\cdot)$ converges if the following 'impatience condition' holds:

$$\sigma^{-1}(r-d) + \left(\frac{\sigma}{2}\right)\sigma_{\ln n}^2 < g - \sigma_{\ln n}^2/2. \quad (1.19)$$

To interpret (1.19), note that if there is no income uncertainty, so that $\sigma_{\ln N}^2 = 0$, then this condition simplifies to $\sigma^{-1}(r-d) < g$, which is the condition in a Modigliani life cycle model with no uncertainty for consumption to grow more slowly than income over life. This in turn means that, since the present value of lifetime consumption and of lifetime income must be equal, initial consumption exceeds initial income; in other words, consumers are impatient, and choose to boost current consumption by borrowing against future income growth. Thus, (1.19) indicates impatience in the no uncertainty case; the $\sigma_{\ln N}^2$ terms in (1.19) reflect the impact of income uncertainty on planned consumption growth (left-hand side) and on income growth (right-hand side). Note that (1.19) can hold even if $r = d$; in practice, however, Carroll focuses in much of his work on the case in which $d > r$.




Figure 1.3 Expected consumption growth and the ratio of cash to labour income

The implications can be seen in figure 1.3, where the curve YY plots expected consumption growth from (1.17) against the ratio of cash on hand to normal labour income. The horizontal lines $\sigma^{-1}(r-d)$ and g indicate, respectively, planned consumption growth in a CEQ Modigliani model, and normal labour income growth. The vertical line $(x/p)^*$ indicates the target ratio of cash on hand to normal labour income, such that if $(x/p)^*$ is attained, households thereafter seek to maintain it constant. $(x/p)^*$ is stable in Carroll's model; thus, the arrows on YY indicate the direction of planned movement.

Figure 1.3 and (1.17) can best be understood together, (1.17) indicates that planned consumption growth varies positively with the variance of future consumption growth. The size of the latter is partly determined by the variance of future labour income, but it is also determined by the quantitative importance of future labour income relative to financial wealth: since financial wealth is not subject to uncertainty, the larger the stock of such wealth relative to uncertain labour income, the greater its ability to cushion consumption from labour income shocks, and hence the smaller the expected variance of future consumption growth. Thus, the variance of future consumption growth is an endogenous variable because it depends on the household's saving and wealth accumulation decisions. Figure 1.3 and (1.17) both indicate that when wealth is low relative to normal income, households seek to accumulate wealth – planned consumption growth is high, indicating that the current consumption level is low – in order to increase their future cushioning ability. Thus, Carroll-type models are

often labelled *buffer stock models*: wealth acts as a buffer against income shocks.

3.4 Impatience versus prudence

As noted above, Carroll mostly focuses on cases in which $d > r$, and invariably $g > 0$ (i.e. growth in normal income). In a CEQ model this parameter combination would generate initial borrowing against future income, and subsequently falling consumption over the lifetime. Why does this not happen in Carroll's framework? Here the specification of the income process is critical. In the model of (1.10)–(1.15) there is a positive probability, q , of zero income in any individual period, so that in a remaining lifetime of N periods, there is a positive probability q^N that income will be zero in all N future periods. Thus, if the household chose to accumulate no buffer stock of financial wealth – or *a fortiori* if it chose to borrow (i.e. negative wealth) – it would face the probability q^N of complete destitution (i.e. $c = 0$) for its remaining lifetime. Because CRRA utility implies infinite marginal utility at $c = 0$, the household will always choose to hold some positive wealth in order to insure against the possibility of destitution. Note the importance here of the assumption, emphasised earlier, that self-insurance is the only option available for countering this threat. I return to this later when the role of social safety nets is considered.

The household's impatience is at war with its prudence. The impatient strand in preferences provides a temptation to overspend now; the prudent strand focuses on the future consequences of overspending in the event of a sequence of extremely unlucky income outcomes. The result is a compromise: the household chooses to accumulate some – but not much – wealth. This produces a now standard finding in the buffer stock literature. Looking at the pattern of life cycle behaviour, saving and wealth accumulation decisions over working life fall into two distinct phases. For typical households up to around age 45–50, buffering behaviour dominates: they choose to accumulate small amounts of wealth as a cushion against income shocks. For typical households from age 45–50 up to retirement, more traditional life cycle motives – saving for retirement – dominate, leading to the accumulation of significantly larger amounts of assets. Thus, all saving choices are ultimately driven by consumption smoothing motives, but for younger working age households it is high frequency smoothing which matters, while for working age households approaching retirement low frequency smoothing is more important. Similar results have been obtained in other recent work using similar frameworks (e.g. Gourinchas and Parker

2001, 2002; Gourinchas 2000). This same argument also generates the ‘consumption tracks income’ phenomenon: once a household has attained its target stock of financial wealth relative to normal income, further growth in normal income leads to approximately equal proportionate growth in consumption, allowing the wealth–income ratio to remain at its target level. Similarly, the CEQ proposition that households consume the permanent income value of any change in either present or future income does not hold. In general, a change in expected future income has little impact on current consumption in the buffer stock model, since it does not affect the ratio x/p which is what drives behaviour. Conversely, a change in current income (and thus in x) can have a large impact on current consumption, depending on the value of x/p relative to the target value.

3.5 *The HSZ framework*

The foregoing discussion focused on Carroll’s (1997) framework. A partly similar, and partly contrasting, framework which has also been influential in the SDP literature is that of Hubbard, Skinner and Zeldes (HSZ) (1994, 1995). The HSZ model differs from Carroll’s in the following ways: (1) As well as labour income uncertainty, HSZ also incorporate uncertainty about the remaining length of life, and about the possibility of medical spending in the event of future ill health; (2) HSZ also incorporate social security, the effect of which is to introduce a consumption floor, $c(\min)$, which is underwritten by the government; i.e. if the household experiences income draws which are sufficiently unlucky that its total available resources fall below the amount needed to consume at $c(\min)$, then the government provides transfer income sufficient to meet the difference; (3) HSZ focus on different baseline parameter combinations from Carroll; in particular, their baseline case sets $r = d$ whereas Carroll assumes $r < d$, and HSZ also assume significantly slower growth of labour income than in Carroll’s model; (4) HSZ assume the existence of borrowing constraints, such that the household’s total wealth can never be negative, i.e.

$$F(t) \geq 0 \quad \text{all } t. \quad (1.20)$$

HSZ simulate lifetime consumption and wealth accumulation profiles for households with different levels of education, and therefore different absolute levels, and lifetime patterns, of labour income. They show that all groups have hump-shaped wealth profiles, but that the presence of the consumption floor has a proportionately large effect on households with

low education attainments, and therefore relatively low lifetime income. For these households, the incentive to accumulate wealth is much reduced – and in some cases disappears altogether – because the guaranteed consumption floor $c(\min)$ is a relatively high proportion of their normal annual income. Thus, for such low-income households it is sensible to save little or nothing, relying on the social security safety net in the event of unfavourable contingencies. By contrast, for households with higher education attainments and therefore higher income levels, it is not attractive to rely on social security because the consumption floor $c(\min)$ is very low relative to their normal income levels, and would therefore involve too large a drop in consumption. Thus, private precautionary saving is relatively more important for high-education than for low-education households. It follows that the presence of social security is a key part of the HSZ model. Of the other differences between the HSZ model and the Carroll framework, the presence of medical uncertainty and length of life uncertainty turns out to have quantitatively small effects, as does the existence of borrowing constraints for some, though not all, parameter combinations.

The final difference between the HSZ and Carroll frameworks concerns the choice of baseline values for the parameters which govern ‘impatience’, i.e. the choice of the time preference rate, d , relative to the interest rate, and also the assumed rate of lifetime income growth. As noted above, compared with Carroll, HSZ assume lower values both for d , and for average lifetime income growth; thus, they assume significantly lower degrees of impatience. Having simulated their model with their own preferred parameter values, HSZ then rework it using typical Carroll-type values; and in his own work Carroll returns the compliment by using HSZ-type parameter values in his framework. What emerges from these exercises is that (surely unsurprisingly) neither framework can easily be made consistent with all features of observed behaviour. For example, (1) HSZ argue (1994: 86) that using Carroll-type parameter values in their framework generates counterfactually small levels of pre-retirement wealth accumulation (because greater impatience reduces the willingness to accumulate pre-retirement savings to supplement guaranteed pension income); but conversely (2) Carroll and Samwick (1997: section 5) argue that using HSZ-type parameter values in their framework causes the sensitivity of the level of accumulated wealth to changes in the degree of uncertainty of future labour income to be counterfactually high (because lower time preference means that unfavourable contingencies relatively far into the future still exert a significant impact on current behaviour).

3.6 Discussion

There are other difficulties with both the Carroll and the HSZ model. The most important concerns the validity, or otherwise, of the entire SDP-type approach which is common to both; I consider this in section 1.4 below. Here, I consider further aspects of the two sets of models within the SDP approach.

Carroll's model hinges critically upon his assumption that at any given time there is a positive probability of zero income over the whole of the remaining lifetime. If this assumption is correct, any household with zero net wealth risks being left destitute for the rest of life, and all households choose to keep assets positive at all times to insure against such an outcome. Because they are also impatient, however, households keep only small amounts of precautionary wealth, at least until near to retirement. This is the mechanism by which Carroll explains the observation that many households hold small positive wealth stocks. Carroll justifies it as follows. He observes that a small fraction of households has zero incomes in any one year in the PSID data, and *assumes* that there is no social security safety net which would catch such households in the event that they were to suffer repeated zero incomes over a run of years. He provides no direct evidence in support of this assumption. Thus, his entire model is driven by an extreme contingency for which no direct supporting evidence is available (Deaton 1992: 192; Pemberton 1998). One could just as easily assume the opposite, i.e. that although events can cause some households *temporarily* to have zero income, any set of circumstances which threatened to cause *prolonged* zero incomes would be caught by a safety net.

Suppose that one assumes the existence of a safety net such that the household can obtain a fraction $Q > 0$ of its 'normal' income in the event of any persistently unfavourable contingency (continued unemployment, chronic ill health, etc.), while otherwise using the Carroll framework unchanged (Q is analogous to HSZ's $c(\min)$). How would this change the model's predictions? We can see at least an approximate answer to this easily, without formally simulating and solving the full model. In Carroll's setup, $Q = 0$ and the typical household in the first part of its life cycle chooses to hold small (near zero) levels of wealth. Thus, in the event that the lifetime destitution contingency were to materialise, this household would be able to consume at fractionally above zero in each year thereafter. Now let Q be positive. Normalise current normal income equal to unity, and, following Carroll, let the annual growth rate of normal income be g and the annual interest rate be r . Taking an infinite horizon for simplicity, the present value

of normal income is:

$$P = (1 + r)/(r - g). \quad (1.21)$$

The present value of the safety net is simply QP .³ Since households are impatient by assumption, they want to spend as much as possible immediately, consistent with being able to consume at just above zero if disaster strikes. Thus, they wish to borrow an amount just short of QP . Since current income is normalised at unity, the ratio of desired borrowing to current income is (slightly less than) QP . Taking typically assumed values for r and g of 0.04 and 0.02, respectively, and taking a safety net value of $Q = 0.2$, implies that the typical household will choose to borrow 10.4 times its current annual income. Thus, a household with current annual income of 40,000 Euros will choose to go into debt to the tune of 416,000 Euros!⁴ Modifying the parameter values or the framework (e.g. finite rather than infinite life) alters the numbers, but so long as Q is positive the inevitable outcome is a large amount of desired borrowing. Pemberton (1998) provides a more formal analysis of this, albeit within the context of a bounded rationality model (see section 1.4) rather than an SDP model, and shows that with Carroll-type parameter assumptions, typical households choose to build up huge amounts of indebtedness over the first few periods of the life cycle, and then pay them off gradually over the rest of life.

Thus, Carroll's model delivers plausible predictions for wealth accumulation only if it assumes $Q = 0$, i.e. no safety net. If instead $Q > 0$, the model delivers non-sensical results, predicting absurdly counterfactual levels of desired borrowing. Since safety nets have been a widespread feature of post-1945 societies in much of Western Europe and North America, it is not clear that the model provides a robust explanation of observed behaviour. Pemberton (1998) argues that this applies more generally, and that even if one assumes much less household impatience than implied by Carroll's benchmark parameter values, households will still desire to accumulate implausibly large amounts of debt.

There are two ways in which the model can be rescued. The first is to assume borrowing constraints, as in HSZ and also as in Deaton (1991, 1992) who combines such constraints with a model otherwise similar to the Carroll framework. Then, no matter how great is the household's impatience, it is unable to incur any net indebtedness. However, there must be some doubt about the underlying plausibility of such a framework. Do

³ Q grows at an annual rate of g , so as to keep in line with normal income.

⁴ To be precise, slightly less than 416,000 Euros.

we really believe that households would ideally choose to go into debt to the tune of many times their current annual income, and are prevented from doing so only by binding borrowing constraints?

An alternative way to rescue the model is to modify the utility function. Suppose that, instead of (1.7), the utility function is as follows:

$$u(c(t)) = \begin{cases} \frac{[c(t) - c_0(t)]^{1-\sigma} - 1}{1 - \sigma} & (\sigma \neq 1) \\ \ln[c(t) - c_0(t)] & (\sigma = 1). \end{cases} \quad (1.22)$$

With preferences as in (1.22), marginal utility approaches infinity as $c(t)$ approaches $c_0(t)$, rather than zero. $c_0(t)$ can be interpreted as a ‘minimum acceptable’ level of consumption, which in the context of low-income countries might be the subsistence level required for physical survival, and in richer countries might instead carry ‘keeping up with the Joneses’ connotations. Suppose now that we equate $c_0(t)$ with $Q(t)$. Then all Carroll’s results still go through, because households always choose to accumulate at least small amounts of wealth, no matter how large is $Q(t)$. This is the approach taken by Zeldes (1989). There are still problems, however, stemming from the requirement that the model should be able to explain observed wealth holdings across different countries and time periods if it is to be convincing. The difficulty is that the value of $Q(t)$ reflects policy decisions made by national governments, and as such it varies across both countries and time periods. Some countries at some times have provided much more generous safety nets than other countries at other times, where ‘generosity’ is measured in the context by the proportion of one’s ‘normal’ earnings which can be replaced by social security if unfavourable contingencies occur. Thus, unless the value of $c_0(t)$ happens to vary across countries and time periods precisely in line with the corresponding variations in $Q(t)$ – unlikely, given that Q is policy-driven whereas c_0 reflects underlying consumer preferences – the model cannot provide a universal explanation of behaviour.

A similar point applies to the HSZ explanation of differences in saving behaviour as between high-education and low-education households. As seen above, this also rests on the role of the safety net, which has a proportionately much larger impact on low-education (and thus relatively low-income) households than on high-education (high-income) households. An implication is that there should be no significant differences in saving and wealth accumulation as between different education and income groups in contexts in which there is no safety net: in such contexts, *ceteris paribus*, the ratio of wealth to permanent income should be independent of the level of permanent income. James, Palumbo and Thomas (JPT)