1

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 intergenerational 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

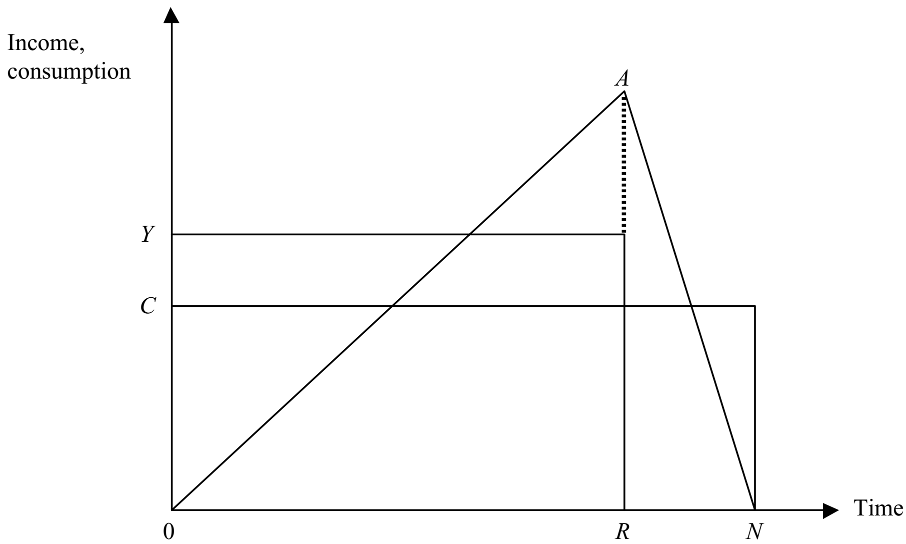

Figure 1.1 The 'Modigliani' diagram

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:

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

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

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

$V(t)$ 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

Cambridge University Press
978-0-521-53403-1 - Dynamic Macroeconomic Analysis: Theory and Policy in General Equilibrium
Edited by Sumru Altug, Jagjit S. Chadha and Charles Nolan
Excerpt
More information

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:

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

[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.

Cambridge University Press
978-0-521-53403-1 - Dynamic Macroeconomic Analysis: Theory and Policy in General Equilibrium
Edited by Sumru Altug, Jagjit S. Chadha and Charles Nolan
Excerpt
More information

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.

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

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

$$E_t c(t+1) = a_0 + a_1 c(t) + e(t+1). \tag{1.6}$$

When $r = d$, $a_0 = 0$ and $a_1 = 1$; a_0 is increasing, and a_1 is decreasing, in $(r - d)$. $e(.)$ 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 $e(.) = 0$ 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:

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

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:

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

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 $E_1c(2)$ denotes mean expected second period consumption, viewed from the vantage

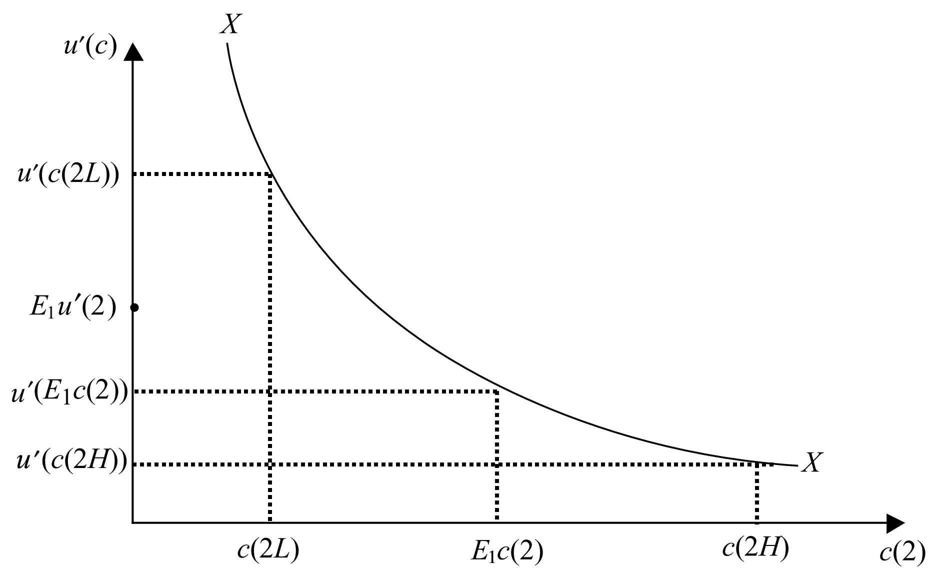

Figure 1.2 Second-period marginal utility of consumption

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

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

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