Chapter 17
Consumption, Income and Wealth

Peter Birch Sørensen and Hans Jørgen Whitta-Jacobsen

2. maj 2003

Private consumption is by far the largest component of aggregate demand for goods and services. Although Chapter 15 showed that consumption is less volatile than investment, changes in the propensity to consume are often the dominant source of changes in total demand, simply because private consumption is typically around three times as large as private investment. A theory of private consumption is therefore an essential building block in any theory of aggregate demand.

Studying private consumption will not only help us to build a short-run model of the business cycle. Since consumption is a basic determinant of economic welfare, our theory of the link between consumption and other macroeconomic variables will also help us understand how business cycles and economic policies affect consumer welfare. Moreover, our theory of consumption will imply a theory of saving, because saving is equal to income minus consumption. Since saving is the basis for capital accumulation, our analysis of consumption is also relevant for the theory of economic growth presented in Book One.

In the section on housing demand in the previous chapter we studied how the consumer allocates his total consumption between housing consumption and consumption of non-durables *within a given time period*. This chapter complements the previous one by analyzing how the consumer will wish to allocate his total consumption *over time*. Since we now wish to explain *aggregate* consumption, this chapter will not elaborate on the previous chapter's analysis of the *composition* of consumption. Thus we will treat consumption as a single aggregate which of course must be thought of as a bundle of commodities, including housing services.

We will start this chapter by briefly restating and discussing the simple Keynesian theory of private consumption. We will then introduce a richer model of consumption to illustrate how consumption is linked to income, wealth and interest rates. In the final part of the chapter we will show how our theory of consumption can be used to analyze the effects of the government's tax and debt policies on aggregate demand.[1]

1 The Consumption Function

The Simple Keynesian Consumption Function

In his famous *General Theory of Employment, Interest and Money* published in 1936, John Maynard Keynes wrote that "... the propensity to consume is a fairly stable function so that, as a rule, the amount of aggregate consumption mainly depends on aggregate income ...". In other words, Keynes argued that real private consumption during period t, denoted by C_t, is mainly determined by real disposable income Y_t^d during that period. In formal terms, $C_t = C\left(Y_t^d\right)$. Keynes went on to argue that: "The fundamental psychological law, upon which we are entitled to depend with great confidence both *a priori* from our knowledge of human nature and from the detailed facts of experience, is that men are disposed, as a rule and on the average, to increase their consumption as income increases, but not by as much as the increase in their income". The claim made by Keynes in this passage is that the *marginal* propensity to consume, $C' \equiv dC_t/dY_t^d$, is positive but less than

[1] This chapter borrows heavily from the teaching note by Henrik Jensen: "Mikrofundament for konventionelle makro-adfærdsrelationer", Københavns Universitets Økonomiske Institut, Marts 1996. We are grateful to Henrik for the inspiration, but of course he should not be held responsible for any shortcomings of our exposition.

one. In a subsequent passage he asserted that "... it is also obvious that a higher absolute level of income will tend, as a rule, to widen the gap between income and consumption. For the satisfaction of the immediate primary needs of a man and his family is usually a stronger motive than the motives towards accumulation, which only acquire effective sway when a margin of comfort has been attained. These reasons will lead, as a rule, to a *greater proportion* of income being saved as real income increases".[2] Thus Keynes believed that the *average* propensity to consume, C_t/Y_t^d, will decrease with the level of income. In other words, the rich are assumed to have a higher average savings rate than the poor.

A consumption function with all of these Keynesian properties is the simple linear one

$$C_t = a + bY_t^d, \qquad a > 0, \qquad 0 < b < 1. \tag{1}$$

In this consumption function the marginal propensity to consume is the constant b, which is below one, and the average propensity to consume is $C_t/Y_t^d = b + a/Y_t^d$ which is obviously decreasing with income.

The consumption function (1) has the virtue of being simple, but there are at least two problems with it. The first problem is theoretical: although it seems plausible that consumption is positively related to current income, it is not clear why the current consumption of an optimizing consumer should depend *only* on current income and nothing else. In short, we would like to know whether a consumption function like (1) is consistent with optimizing behaviour.

The second problem is empirical: although microeconomic cross-section data on the relationship between consumption and income for different families within a given period do indicate that the rich save a larger fraction of their current income than the poor,

[2] This quote and the two previous ones can be found on pp. 96-97 of John Maynard Keynes: *The General Theory of Employment, Interest and Money*, Macmillan Press, London and Basingstoke, 1936. These brief quotes do not do full justice to Keynes' theory of consumption. He did in fact discuss a host of other factors likely to influence private consumption. Still, it is fair to say that as a first approximation, he believed that current consumption depends mainly on current income.

macroeconomic time series data for most countries indicate that the ratio of aggregate consumption to aggregate income is roughly constant over the long run. These apparently contradictory stylized facts are illustrated in Figures 17.1.a and 17.1.b.

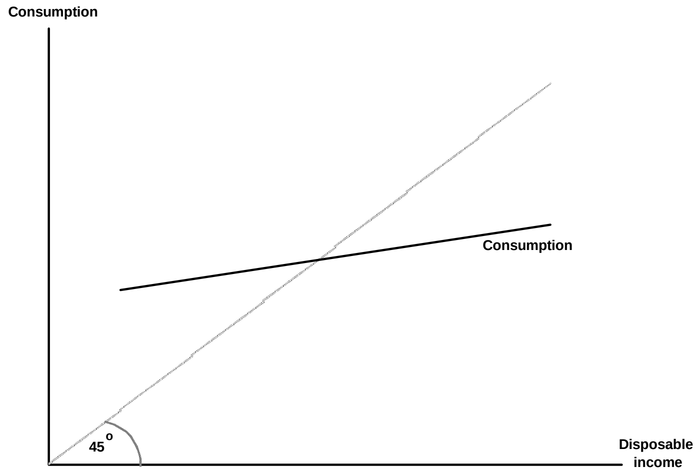

Figure 17.1a: Stylized relationship between income and consumption in microeconomic cross-section data

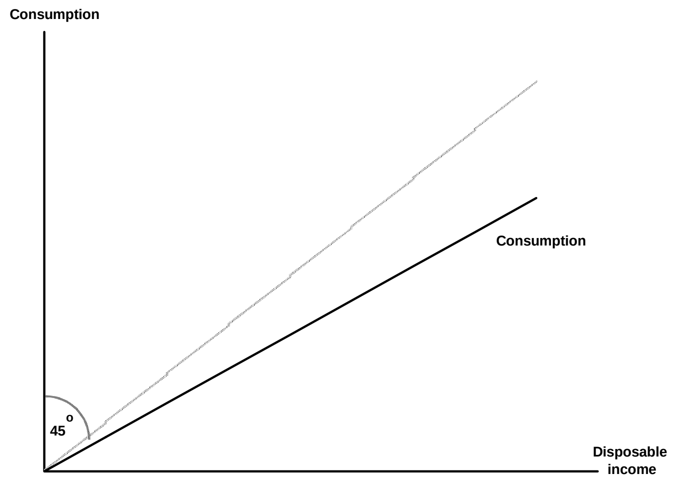

Figure 17.1b: Stylized relationship between income and consumption in macroeconomic time series data

The rough long-run constancy of the average propensity to consume was used extensively in the growth theory of Book One through the simple consumption and savings behaviour assumed there, and it was documented in Figure 2.2 of Chapter 2 for six different countries. It is illustrated again for the USA and Denmark in Figure 17.2. Apart from a temporary drop due to consumption rationing during World War II, the average propensity to consume in the U.S. has been remarkably stable over the long run, despite the tremendous growth in income since 1929. The Danish propensity to consume has been more volatile, but without any systematic trend. Figure 17.2 clearly contradicts equation (1) which implies that the ratio of consumption to income should *decline over time* as income grows. Instead of equation (1), we therefore need a theory of consumption which has a solid microtheoretic foundation and which is able to explain why we observe different relationships between consumption and income in microeconomic cross-section data and in macroeconomic time series data. In the rest of this chapter we shall try to build such a theory.

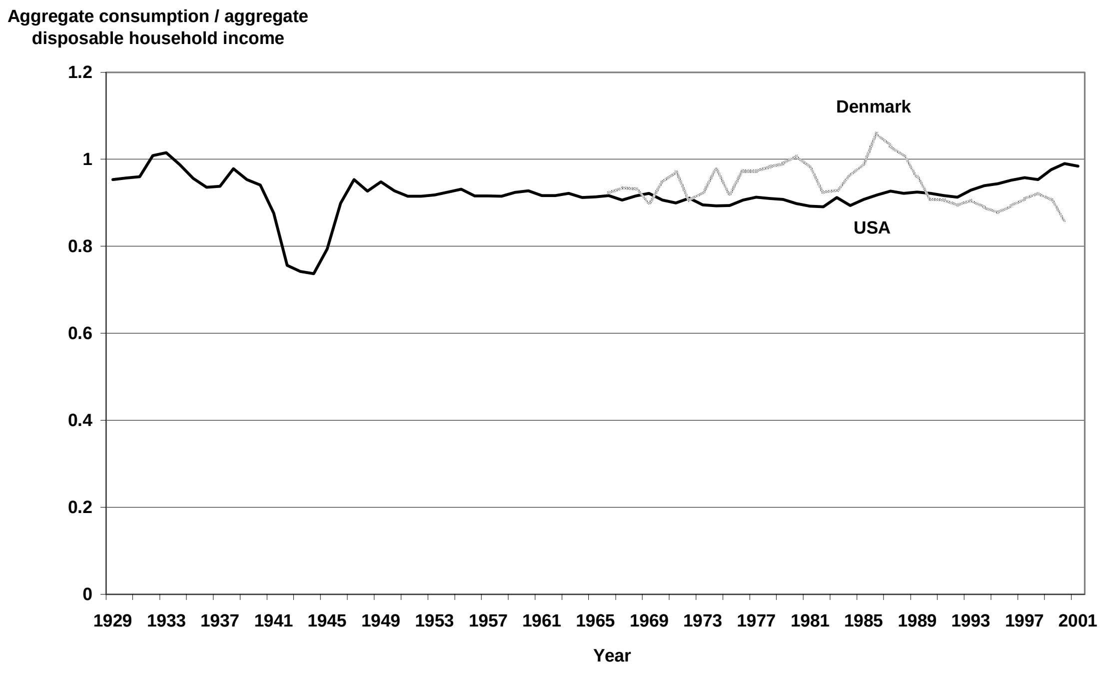

Figure 17.2: The average propensity to consume in USA and Denmark

Source: National Income Accounts, Bureau of Economic Analysis and ADAM database, Statistics Denmark

Consumer Preferences

The starting point for a micro-based theory of consumption is a specification of consumer preferences. Consider a consumer who plans for a certain finite time horizon. We will divide this time interval into two periods which may be thought of as 'the present' (the current period 1) and 'the future' (period 2). The limitation to only two periods is just a simplification; one can show that all our qualitative conclusions will continue to hold in a setting with many periods.

In each period t $(t = 1, 2)$ the consumer derives utility $u(C_t)$ from consumption. However, because the consumer is 'impatient', he prefers a unit of utility today to a unit of utility tomorrow. When evaluated at the beginning of period 1, the consumer's *lifetime* utility U is therefore given by

$$U = u(C_1) + \frac{u(C_2)}{1+\phi}, \qquad u' > 0, \qquad u'' < 0, \qquad \phi > 0. \tag{2}$$

The consumer's impatience is captured by the parameter ϕ which is referred to as *the rate of time preference*. The positive rate of time preference means that a given amount of consumption today is valued more highly than a similar amount of consumption tomorrow. On the other hand, as long as ϕ is not infinitely high, the consumer is not indifferent about the future, and he will then have to decide how to allocate his consumption optimally over time. The assumptions $u' > 0$ and $u'' < 0$ reflect that the marginal utility of consumption in any period is positive but decreasing. An increase in consumption in any period will thus reduce the marginal utility gain from a further consumption increase in that period.

Our theory of consumption will be derived from the assumption that the consumer trades off present against future consumption so as to maximize the lifetime utility function (2). The terms of this trade-off will depend on the consumer's intertemporal budget constraint to which we now turn.

The Intertemporal Budget Constraint

When specifying the consumer's budget constraints for the two periods, we will assume that *capital markets are perfect.* This means that the consumer can freely lend and borrow as much as he likes at the going market rate of interest. In practice some consumers may face *credit constraints* preventing them from borrowing as much as they would have preferred at the going interest rate. We will discuss credit constraints later in the chapter, and Exercise 1 asks you to consider their implications in detail, but for the moment we will stick to the assumption of perfect capital markets.

We specify the consumer's budget constraints in real terms. At the beginning of period 1 the consumer is endowed with a predetermined stock of real financial wealth V_1. During period 1 he earns real labour income Y_1^L, pays the real amount of taxes T_1, and spends the real amount C_1 on consumption. For convenience we assume that all payments are made at the beginning of the period.[3] After having received his income and incurred his expenditure on taxes and consumption, the consumer then has an amount $V_1+Y_1^L-T_1-C_1$ left over for investment in interest-bearing financial assets. If the real interest rate is r, the consumer will therefore end up with a real stock of financial wealth $V_2 = (1+r)\left(V_1 + Y_1^L - T_1 - C_1\right)$ at the beginning of period 2. Hence the budget constraint for period 1 is

$$V_2 = (1+r)\left(V_1 + Y_1^L - T_1 - C_1\right). \tag{3}$$

Note that V_2 may well be negative. In that case the consumer is a net borrower during period 1. Since he does not plan any consumption beyond period 2, he will simply spend all his resources during that period, including the financial wealth he accumulated during

[3] If some or all payments were instead made at the end of the period (as we assumed in the housing market model in the previous chapter), we would obtain a consumption function with the same qualitative properties as those described below. However, the assumption that payments take place at the start of each period leads to slightly more elegant analytical expressions. Whenever we divide the time axis into discrete finite intervals, there is no objectively 'correct' assumption on the timing of payments (beginning-of-period versus end-of-period). Hence we are free to choose that assumption on timing which is most convenient for analytical purposes.

period 1. Using the same notation as before, we may therefore write the budget constraint for period 2 as

$$C_2 = V_2 + Y_2^L - T_2. \tag{4}$$

Equation (4) states that the consumer will spend his initial financial wealth plus his after-tax labour income on consumption during period 2. Notice that if he has borrowed during period 1 so that $V_2 < 0$, he must reserve part of his period 2 labour income for repayment of his debt.

It will turn out to be convenient to consolidate (3) and (4) into a single constraint. We therefore use (3) to eliminate V_2 from (4), divide through by $(1+r)$ on both sides of the resulting equation and rearrange to get

$$C_1 + \frac{C_2}{1+r} = V_1 + Y_1^L - T_1 + \frac{Y_2^L - T_2}{1+r}. \tag{5}$$

Equation (5) is the consumer's *intertemporal budget constraint*. It states that the present value of the consumer' lifetime consumption (the left-hand side) must equal the present value of his after-tax labour income plus his initial financial wealth (the right-hand side).[4] In other words, with a perfect capital market current consumption does not have to equal current income, but over the life cycle the consumer cannot spend any more than his total resources. These resources consist of his labour income and his initial financial wealth.

We can write (5) in a simpler form by introducing

$$H_1 \equiv Y_1^L - T_1 + \frac{Y_2^L - T_2}{1+r}. \tag{6}$$

The variable H_1 is the present value of the consumer's disposable lifetime labour income, and it is referred to as his *human wealth* or *human capital* because it measures his capitalized earnings potential in the labour market. Note that H_1 carries the time subscript

[4] Of course the consumer could choose to consume less than his total lifetime resources, but since an increase in consumption today or tomorrow will always increase his lifetime utility, he will always choose to consume as much as his budget constraint permits. This is why (5) is written with an equality sign rather than an inequality sign.

1 because it includes labour income from period 1 and onwards. Inserting (6) into (5) we get

$$C_1 + \frac{C_2}{1+r} = V_1 + H_1. \tag{7}$$

Hence the consumer's intertemporal budget constraint simply states that the present value of his real lifetime consumption is constrained by his *total real initial wealth*, consisting of the sum of his financial wealth and his human wealth.

We shall now study how consumers will actually wish to allocate consumption over time.

The Allocation of Consumption over Time

The consumer chooses his time path of consumption so as to maximize his lifetime utility function (2) subject to his intertemporal budget constraint (7). The consumer takes his wage rate as given, and we assume that his working hours are institutionally determined, say, by collective bargaining agreements or by law. This means that the labour incomes Y_1^L and Y_2^L and thereby H_1 are exogenously given to the consumer.[5] Since the market value of financial assets is also beyond the control of the individual consumer, it follows that he takes his total initial wealth $V_1 + H_1$ as given when optimizing his consumption. Because we are mainly interested in studying the determinants of current consumption C_1, we will use the intertemporal budget constraint (7) to eliminate C_2 from the lifetime utility function (2). Doing that, we obtain

$$U = u(C_1) + \frac{u\left((1+r)(V_1 + H_1 - C_1)\right)}{1+\phi}. \tag{8}$$

The consumer's problem then boils down to choosing C_1 so as to maximize (8). The first-order condition for a maximum is: $dU/dC_1 = u'(C_1) - \frac{1+r}{1+\phi}u'\left((1+r)(V_1 + H_1 - C_1)\right) = 0$,

[5] In Exercise 2 you will be asked to study the determination of consumption in the more complicated case where the consumer can choose his working hours and hence his labour income.

which we can write as

$$u'(C_1) = \frac{1+r}{1+\phi}u'(C_2), \tag{9}$$

or as

$$\frac{u'(C_1)}{u'(C_2)/(1+\phi)} \equiv MRS(C_2 : C_1) = 1 + r. \tag{10}$$

Equations (9) and (10) are just two alternative ways of writing the consumer's optimum condition. Consider first the interpretation of (9). If the consumer increases consumption by one unit in period 1, his lifetime utility will increase by $u'(C_1)$. If he chooses instead to save an extra unit and invest the funds in the capital market to earn the real interest rate r, he will be able to increase his consumption in period 2 by $1+r$ units. This will generate an increase in lifetime utility equal to $\frac{1+r}{1+\phi}u'(C_2)$. According to (9) these two alternatives must be equally attractive. In other words, in optimum the consumer must be indifferent between consuming an extra unit today and saving an extra unit today.

Equation (10) is a version of the usual optimum condition that the consumer's marginal rate of substitution, MRS, between any two goods must equal the price ratio between the two goods. In this case the two goods are 'present consumption' C_1 and 'future consumption' C_2. The optimum condition (10) is illustrated in Figure 17.3.

To interpret the figure, note that along any consumer indifference curve lifetime utility is constant. According to the lifetime utility function (2) a constant utility level implies $dU = u'(C_1)\,dC_1 + \frac{u'(C_2)}{1+\phi}dC_2 = 0$, or

$$-\frac{dC_2}{dC_1} = \frac{u'(C_1)}{u'(C_2)/(1+\phi)}. \tag{11}$$

Equation (11) shows that the numerical slope of the indifference curve is equal to the marginal rate of substitution between present and future consumption, $MRS(C_2 : C_1)$, defined as the ratio of the marginal lifetime utility of present consumption $u'(C_1)$ to the marginal lifetime utility of future consumption $u'(C_2)/(1+\phi)$. Figure 17.3 also shows that the numerical slope of the consumer's lifetime budget constraint equals $1+r$: if the consumer gives up one unit of present consumption, he can have $1 + r$ additional units of future consumption because his saving earns the real interest rate r. Hence we can say that $1 + r$ is the *relative price of present consumption*, since it measures the amount of future consumption which must be given up to enable the consumer to increase present consumption by one unit. When the consumer attains the highest possible level of lifetime utility consistent with his lifetime budget constraint, we see from Figure 2 that the slope of his indifference curve must equal the slope of his budget constraint, given by the relative price of present consumption. This is exactly what equation (10) says.

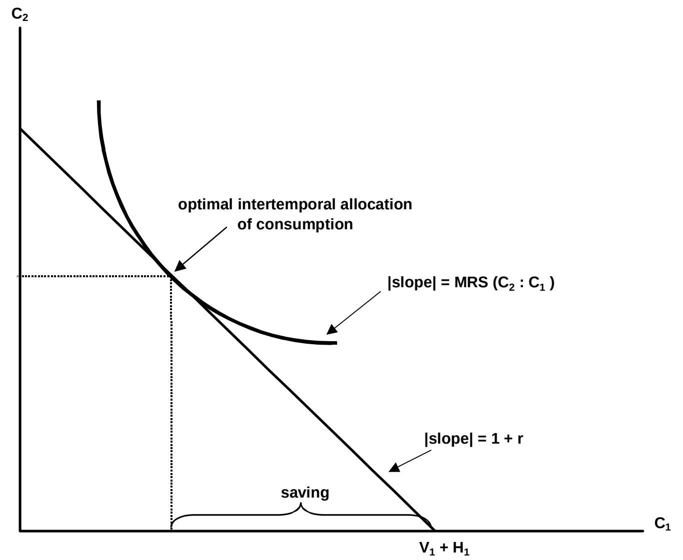

Figure 17.3: The consumer's optimal intertemporal allocation of consumption

The optimum consumption rule (10) has been named the *Keynes-Ramsey rule* after its discoverers. Although formally derived by Cambridge economist Frank Ramsey in 1928,[6] the rule had been verbally anticipated by the very same John Maynard Keynes who later (in his *General Theory* from 1936) came to believe in a consumption function like (1)! As we will discuss later in this chapter (and in Exercise 1), for some households current

[6]Frank Ramsey: "A Mathematical Theory of Saving", *Economic Journal*, vol. 38 (December), pp. 543-559. This was one of two path-breaking articles published by the unusually gifted Frank Ramsey before his premature death at the age of 25.

consumption is inded likely to depend only on current income, as postulated in the consumption function (1). But we shall also see that the consumption function implied by the Keynes-Ramsey rule (10) is consistent with several empirical observations on consumption which are not consistent with the simple Keynesian consumption function.

One important implication of (10) is that consumers will typically want to use the capital market to smooth their consumption over time. This is seen most clearly in the benchmark case where the real interest rate equals the rate of time preference ($r = \phi$) so that the consumer's impatience ϕ is exactly offset by the capital market reward r for postponing consumption. When $r = \phi$ condition (10) can only be met if $C_1 = C_2$, that is if consumption is constant over the consumer's life cycle. Suppose for simplicity that the consumer starts out in period 1 with zero financial wealth, $V_1 = 0$. Unless his labour income happens to be the same in periods 1 and 2, he will then have to engage in financial saving in one period and financial dissaving in the other period to keep consumption constant over time.

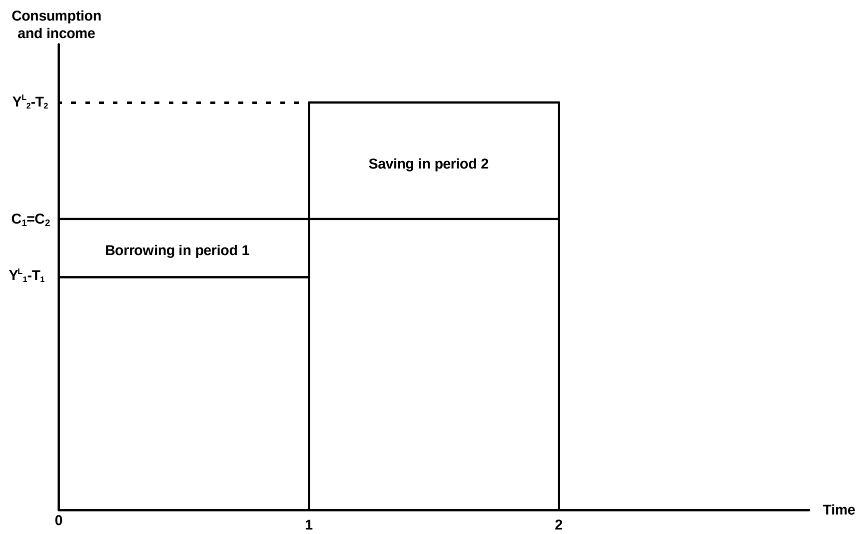

Figure 17.4a: Consumption smoothing for a consumer with relatively low income during period 1 ($\mathbf{V}_1 = 0$)

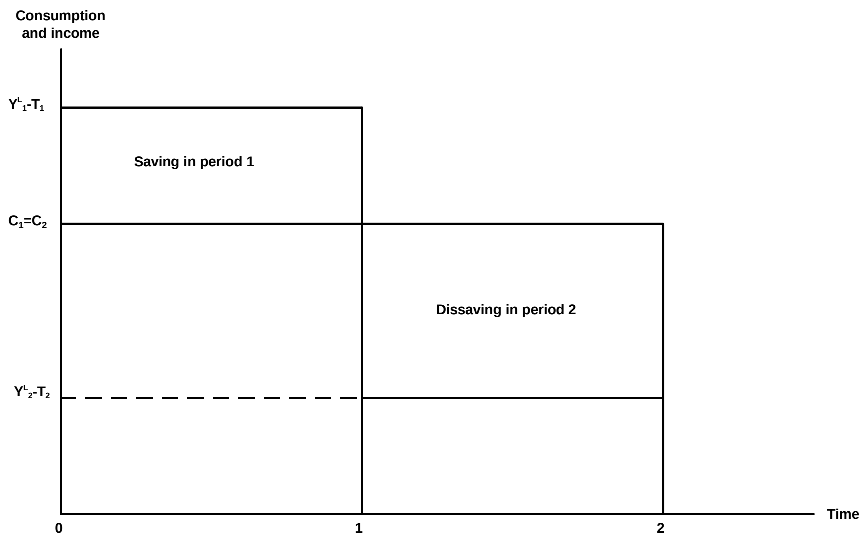

Figure 17.4b: Consumption smoothing for a consumer with relatively high income during period 1 ($\mathbf{V}_1 = 0$)

Figure 17.4.a shows the case where labour income in period 1 is lower than labour income in period 2. The consumer will then want to borrow in period 1 to smooth his consumption over the life cycle even though this means that he has to reserve part of his higher future labour income for payment of interest. Figure 17.4.b shows the alternative case where future labour income is lower than current labour income (perhaps because the consumer plans to retire some time in period 2). In this case the consumer will want to save part of his current labour income and will partly finance his period 2 consumption out of his previous savings.

If there were no capital market allowing saving and dissaving, consumption would have to equal income within each period. If income is low in one period and high in another, the marginal utility of consumption will then differ across periods, because marginal utility decreases with increasing consumption. The consumer will therefore enjoy a welfare gain to the extent that he is able to smooth consumption by borrowing or saving via the capital market. The basic message is that capital markets enable consumers to decouple current consumption from current income, and utility-maximizing consumers will typically want to take advantage of that in order to spread their consumption more evenly over time. Although total lifetime consumption is constrained by total lifetime resources (wealth), we should not necessarily expect to observe a close link between current consumption and current income.

Do consumers actually smooth their consumption over their working life? A glance at the upper two curves in Figure 17.5 might suggest that the answer is 'no'. The figure tracks data for the average consumption-age profile and the average income-age profile for a cohort of British married couples. The figures for income and consumption have been deflated by the consumer price index and transformed into logarithms. The age of the household is identified with the age of the female household member. We see that both consumption and income follow a hump-shaped pattern over the life cycle, peaking a little before the age of 50. In particular, current consumption seems to follow current income fairly closely, in apparent contrast to the hypothesis of consumption-smoothing.

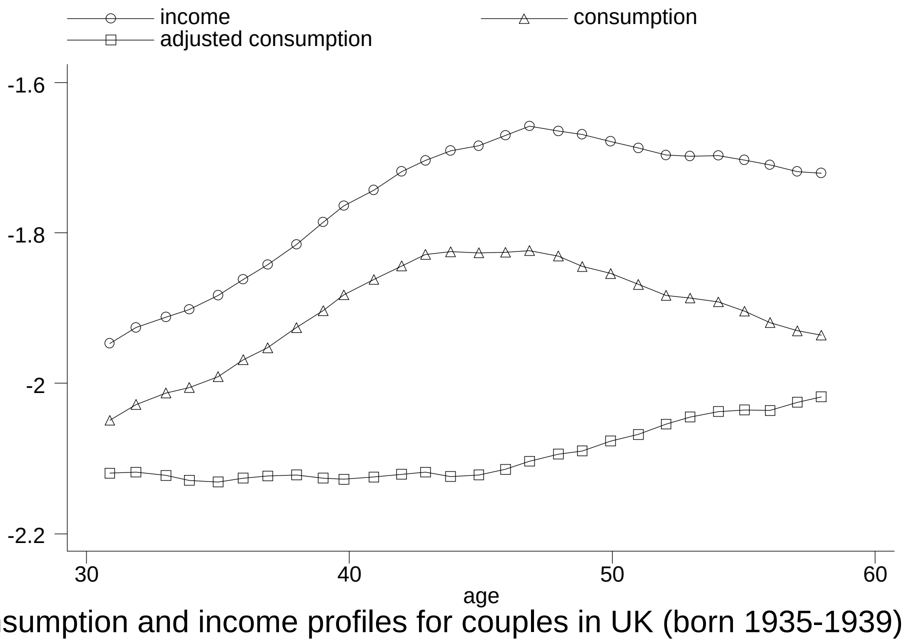

Figure 17.5: Consumption and income profiles for households in the United Kingdom (born 1935-1939)

Source: Data are from the UK Family Expenditure Survey, kindly provided by Mette Ejrnæs.

However, Martin Browning and Mette Ejrnæs of Copenhagen University have demonstrated that if one corrects the consumption figures for the systematic impact of differences in the number of children across households, consumers *do* actually tend to smooth their consumption over time.[7] This is illustrated by the bottom curve in Figure 17.5 which shows 'adjusted consumption', that is, the level of consumption adjusted for the estimated impact of the numbers and ages of children on the consumption needs of the household. We see that the adjusted consumption-age profile is much flatter than the income profile, suggesting that consumers do indeed prefer to smooth consumption per household member, in accordance with our theory.

The Determinants of Current Consumption

The Keynes-Ramsey rule (10) only provides an implicit solution to the consumer's problem. To derive an explicit analytical solution for current consumption C_1, we need to specify the form of the consumer's utility function. One specification which has often been used in economic research is the following one:

$$u(C_t) = \frac{\sigma}{\sigma - 1} C_t^{\frac{\sigma-1}{\sigma}} \qquad \text{for } \sigma > 0, \neq 1, \tag{12.a}$$

$$u(C_t) = \ln C_t \qquad \text{for } \sigma = 1. \tag{12.b}$$

As you may easily verify, this specification satisfies the assumptions in (2) that $u' > 0$ and $u'' < 0$. To interpret the parameter σ in (12.a), we introduce the concept of the *intertemporal elasticity of substitution in consumption*, IES, defined as the percentage change in the ratio of future to present consumption (C_2/C_1) implied by a one percent change in the consumer's marginal rate of substitution, $MRS\,(C_2 : C_1)$:

$$IES \equiv \frac{d\,(C_2/C_1)\,/\,(C_2/C_1)}{dMRS\,(C_2 : C_1)\,/MRS\,(C_2 : C_1)} = \frac{d\ln(C_2/C_1)}{d\ln MRS\,(C_2 : C_1)}. \tag{13}$$

[7] See Martin Browning and Mette Ejrnæs: Consumption and Children, Working Paper, Institute of Economics, University of Copenhagen, February 2002.

The IES measures the degree to which the consumer is willing to substitute future for present consumption. This is illustrated in Figure 17.6 where we recall that $MRS\,(C_2 : C_1)$ measures the numerical slope of the consumer's indifference curve. The slope of the straight line from the origin through point E^o measures the consumption ratio (C_2^o/C_1^o) prevailing when the marginal rate of substitution equals $MRS\,(C_2^o : C_1^o)$. As we move up the indifference curve from point E^o to point E^1, thereby raising the consumption ratio from (C_2^o/C_1^o) to (C_2^1/C_1^1), the consumer becomes less willing to trade present consumption for future consumption. This is reflected by the increase in the marginal rate of substitution, that is, by the fact that the numerical slope of the indifference curve is steeper at point E^1 than at E^o. It is obvious that, for any given increase in $MRS\,(C_2 : C_1)$, the rise in the consumption ratio (C_2/C_1) will be *greater* the *flatter* the indifference curve, that is, the easier it is to substitute future for present consumption. In other words, the more the consumer is willing to engage in intertemporal substitution in consumption, the greater will be the value of IES defined in (13).

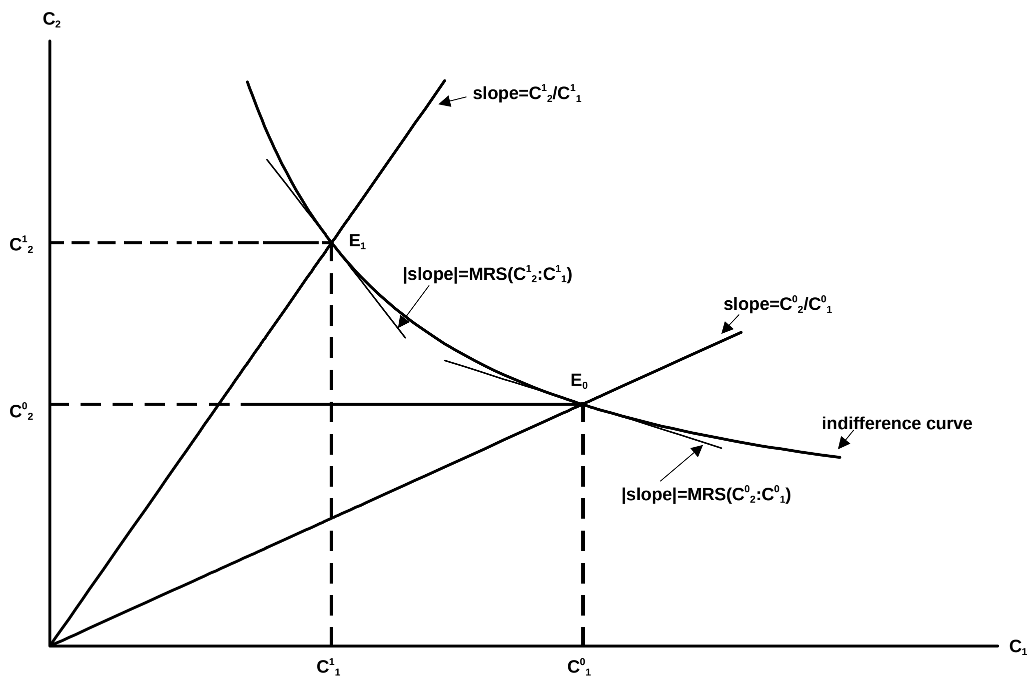

Figure 17.6: The relation between the consumption ratio ($\mathbf{C_2/C_1}$) and the

marginal rate of substitution

Since $u'(C_t) = C_t^{-1/\sigma}$, we may now write the MRS, $\frac{u'(C_1)}{u'(C_2)/(1+\phi)}$, as

$$MRS\,(C_2 : C_1) = (1+\phi)\left(\frac{C_2}{C_1}\right)^{1/\sigma},$$

and hence

$$\ln MRS\,(C_2 : C_1) = \ln(1+\phi) + \frac{1}{\sigma}\ln\left(\frac{C_2}{C_1}\right),$$

from which

$$IES = \frac{d\ln(C_2/C_1)}{d\ln MRS\,(C_2 : C_1)} = \sigma. \tag{14}$$

The utility function (12) thus has the property that the intertemporal elasticity of substitution is constant and equal to σ. By varying σ we may therefore study the effects of variations in the consumer's willingness to substitute consumption over time. As σ tends to zero, the consumer becomes quite unwilling to trade present for future consumption, and his indifference curves become rectangular, as shown in Figure 17.7. On the other hand, as σ tends to infinity, substitution possibilities also become infinite, and the indifference curves converge to straight lines.

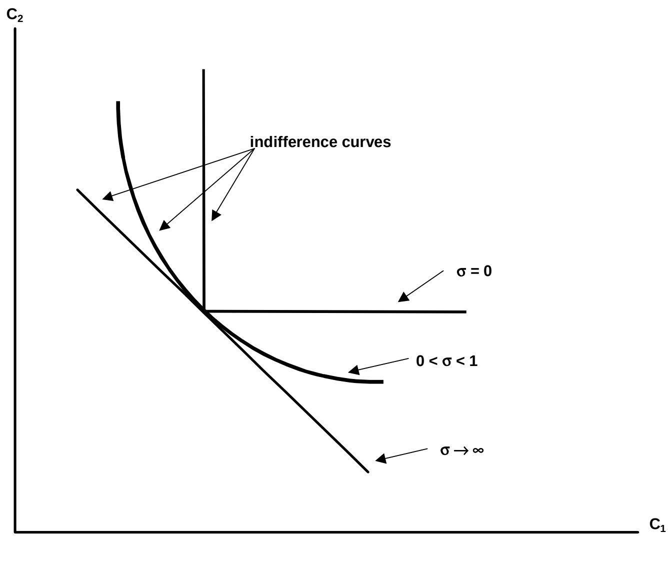

Figure 17.7: The relation between the shape of the indifference curve and the intertemporal substitution elasticity

We are now ready to derive the consumption function. Substituting the expression $(1+\phi)(C_2/C_1)^{1/\sigma}$ for $MRS\,(C_2 : C_1)$ into the condition $MRS(C_2 : C_1) = 1+r$ for an optimal intertemporal allocation of consumption we get

$$C_2 = \left(\frac{1+r}{1+\phi}\right)^{\sigma} C_1. \tag{15}$$

This may be inserted into the lifetime budget constraint $C_1 + \frac{C_2}{1+r} = V_1 + H_1$ to give $C_1 + (1+r)^{\sigma-1}(1+\phi)^{-\sigma} C_1 = V_1 + H_1$, implying

$$C_1 = \theta\,(V_1 + H_1), \qquad 0 < \theta \equiv \frac{1}{1+(1+r)^{\sigma-1}(1+\phi)^{-\sigma}} < 1. \tag{16}$$

Equation (16) is seen to be rather different from the simple Keynesian consumption function (1). According to (16) *current consumption is proportional to current wealth.* The propensity to consume out of current wealth (θ) is positive but less than one, and its magnitude will depend on the real rate of interest (in a manner to be studied in detail below).

2 The Properties of the Consumption Function

Consumption and Income

Let us now explore the implications of (16) for the relationship between current consumption C_1 and current disposable labour income, $Y_1^d \equiv Y_1^L - T_1$. If we insert our definition of human wealth (6) into (16), we get

$$C_1 = \theta\left(Y_1^d + \frac{Y_2^d}{1+r} + V_1\right), \qquad Y_t^d \equiv Y_t^L - T_t, \qquad t = 1, 2. \tag{17}$$

This may be rewritten as

$$C_1 = \widehat{\theta} Y_1^d, \tag{18}$$

$$\widehat{\theta} \equiv \theta\left(1 + \frac{R}{1+r} + v_1\right), \qquad R \equiv \frac{Y_2^d}{Y_1^d}, \qquad v_1 \equiv \frac{V_1}{Y_1^d}, \tag{19}$$

where $\widehat{\theta}$ measures the propensity to consume (out of) current income, R is the ratio of future to current income, and v_1 is the current wealth-income ratio. On the basis of (18) and (19) we may now explain the empirical puzzle that the average propensity to consume current income seems to decrease with income when we consider microeconomic cross-section data whereas it seems to be roughly constant when we consider macroeconomic time series data.

Consider first the observed cross-section relationship between income and consumption suggesting that the rich have a higher average savings rate than the poor. Within any given time period, some of those individuals who record a high current income level cannot expect to earn similar high incomes in the future. One important reason is that workers must some day retire from the labour market in which case they will no longer earn income from labour. Moreover, for selfemployed individuals a high current income will sometimes reflect that business has been unusually good during that year. Thus, if a high current income level is expected to be *temporary*, the ratio of future to current income (R) will be relatively low, and according to (19) this will imply a low average propensity to consume current income. In a similar way, some consumers with low current incomes may have good reasons to expect that they will earn more in the future. This will typically be the case for university students and other young people who have not yet realized their earnings potential in the labour market, and for entrepreneurs who have just started up a business or who experience unusually bad business during the current year. For such individuals the ratio R in (19) will be high, and hence they will have a high propensity to consume out of their low current income. If some of the observed variation in income levels in a cross-section of consumers reflects temporary factors, equation (19) may thus explain why the average propensity to consume appears to fall as income goes up.

In a celebrated article, Franco Modigliani and Richard Brumberg provided the first statement of the so-called *life cycle theory* of consumption according to which consumers in different stages of the life cycle will have different propensities to consume current income because of their desire to smooth consumption over time.[8] In a similarly famous contribution, Milton Friedman developed the so-called *permanent income hypothesis* which says that transitory changes in income will mainly lead to temporary changes in savings, whereas current consumption will depend on the consumer's *permanent* income, that is, his expected long-run average income (which is proportional to his total wealth).[9] Thus these writers pointed out that some of the high incomes observed during a given year are only temporarily high and hence will not induce high levels of current consumption, and some of the low incomes recorded in a given period are just temporarily low and hence will not cause similarly low levels of consumption. Because of this, a cross-section analysis of the relationship between current consumption and current income in any given year will give the impression that high-income people have a lower average propensity to consume than low-income people.

Let us next see how equations (17)-(19) may explain the empirical observation that the average propensity to consume is roughly constant in the long run when we look at aggregate time series data. If we denote the growth rate of real income by g, we have $Y_2^d \equiv (1+g)\,Y_1^d$ which may by inserted into (19) to give

$$\hat{\theta} = 1 + \frac{1+g}{1+r} + v_1. \tag{20}$$

Whereas the growth rate g fluctuates in the short run, on average it has been fairly constant over the long run, as we have seen in Figure 4.9 and Figure 15.1. Moreover, the real interest rate r shows no systematic long-run trend as we saw in Figure 4.10, and over

[8] See Franco Modigliani and Richard Bromberg (1954): "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data". In Kenneth K. Kurihara, ed., *Post-Keynesian Economics*, pp. 388-436. New Brunswick, NJ: Rutgers University Press.

[9] Milton Friedman (1957): *A Theory of the Consumption Function.* Princeton, NJ: Princeton University Press. Note that the consumer's stock of total wealth equals the discounted value of his expected future income. Permanent income is that hypothetical constant level of income which has the same present value as the consumer's expected future income stream.

the long run wealth tends to grow at the same rate as income so that the long-run wealth-income ratio $v_1 \equiv V_1/Y_1^d$ is roughly constant. It then follows from (20) that the long-run average propensity to consume will also tend to be constant, as we do indeed observe.

To sum up, our reconciliation of the microeconomic cross-section data and the macroeconomic time series data runs as follows: a temporary increase in income for an individual consumer will reduce his expected values of the parameters R and g in (19) and (20), and will probably also reduce his short-run wealth-income ratio v_1. For these reasons the average propensity to consume will tend to fall with rising income levels in a cross-section of consumers. Over the long run, the average growth rate of income across all consumers is roughly constant, and wealth moves roughly in line with income. From (20) this implies a constant long-run average propensity to consume at the macro level.

Consumption and Wealth

While the average propensity to consume current income tends to be constant in the long run, it fluctuates quite a lot in the short run, as indicated in Figure 17.8. Equation (20) suggests three possible reasons for this: 1) short-run changes in the expected growth rate of income (g), 2) short-run changes in the wealth-income ratio (v_1), and 3) short-run changes in the real rate of interest. According to (20) a higher expected income growth or a rise in the market value of financial wealth relative to current income will increase the propensity to consume. Note that changes in g and v_1 will often go hand in hand, since a higher expected income growth is likely to drive up the market prices of stocks and owner-occupied housing, thereby increasing V_1, because a higher expected growth rate will tend to raise expected future corporate dividends and to drive up the demand for housing.

Figure 17.8 shows that there is indeed a fairly close empirical relationship between the wealth-income ratio and the average propensity to consume, as our theory of consumption would lead us to expect[10].

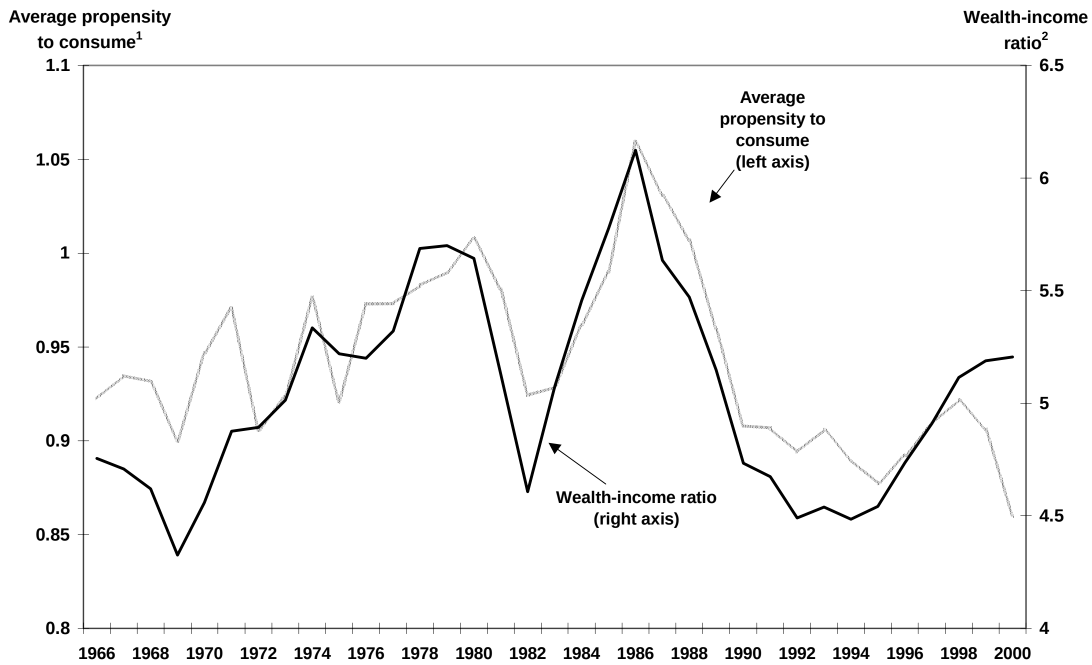

Figure 17.8: Private consumption and wealth in Denmark

[1] Ratio of private consumption to disposable income.

[2] Ratio of private wealth to disposable income.

Source: ADAM database, Statistics Denmark

Consumption and Interest Rates

The third determinant of the propensity to consume current income is the real interest rate r. The interest rate affects consumption via three different channels. First of all, it influences the propensity θ to consume out of wealth. Second, it affects the market value of financial wealth. Third, the real interest rate also affects the value of human wealth. We will consider each of these channels in turn.

[10] It is worth noting that even though he did not include financial wealth as an explanatory variable in his basic consumption function, Keynes was fully aware of the potential influence of wealth on private consumption. On pp. 92-93 of his *General Theory* he wrote: "Windfall changes in capital-values ... are of much more importance in modifying the propensity to consume, since they will bear no stable or regular relationship to the amount of income. The consumption of the wealth-owning class may be extremely susceptible to unforeseen changes in the money-value of its wealth. This should be classified amongst the major factors capable of causing short-period changes in the propensity to consume".

In (16) the propensity to consume wealth was defined as

$$\theta \equiv \frac{1}{1 + (1+r)^{\sigma-1}(1+\phi)^{-\sigma}}. \tag{21}$$

Clearly this propensity increases with r if $\sigma > 1$, and decreases if $\sigma < 1$, and we cannot tell on theoretical ground say which direction is right as σ can be either smaller or larger than one. The indeterminacy in how r affects is due to offsetting substitution and income effects. On the one hand an increase in the real interest rate raises the relative price $1+r$ of current consumption. Ceteris paribus, this will induce the consumer to substitute future for present consumption through an increase in current saving. Obviously this *substitution effect* will tend to reduce the propensity to consume in the current period. At the same time the higher interest rate also increases the amount of future consumption generated by a given amount of current saving. Hence the consumer can afford a higher level of current consumption without having to reduce his future consumption. Other things equal, this *income effect* of the rise in the interest rate works in favour of a rise in current as well as future consumption and thus tends to increase θ.

Since σ measures the strength of the substitution effect, it is not surprising that the sign of the derivative $d\theta/dr$ depends on the magnitude of σ. We see that, for an intertemporal substitution elasticity equal to 1, the income and substitution effects will exactly offset each other, and the propensity to consume will be unaffected. Over the years many researchers have tried to estimate the empirical magnitude of the intertemporal substitution elasticity σ. Most authors have found values of σ well below 1, suggesting that the propensity to consume wealth will *rise* when the real interest rate goes up.

However, a higher interest rate will also influence the level of wealth itself. As already mentioned, the stock of financial wealth V_1 includes the value of stock and of housing capital. In Chapter 16 we saw that both of these important wealth components will be negatively affected by a rise in the real interest rate. Thus, a rise in r means that expected future dividends are discounted more heavily, leading to a fall in share prices. Moreover, by raising the user cost of owner-occupied housing, an increase in r reduces housing demand which in turn drives down the market value of the existing housing stock. For these reasons a rise in the real interest rate will reduce the wealth-income ratio v_1.

In addition, the higher interest rate means that expected future labour income is discounted more heavily. As a consequence, the value of human wealth specified in (6) goes down. Intuitively, a higher real interest rate makes future labour income less valuable by making it easier to attain a given level of future consumption through saving out of current income.

The fall in human and financial wealth induced by a rise in the real interest rate clearly tends to reduce the propensity to consume current income, but since we have seen that there may be an offsetting increase in the propensity to consume wealth (θ), the net effect of a rise in r on the ratio of current consumption to current income ($\widehat{\theta}$) is ambiguous. Against this background it is not surprising that empirical research has found it difficult to document a strong effect of real interest rates on consumption, although the dominant view is that a rise in r tends to reduce consumption because of the negative impact on wealth.

3 Consumption, Taxation and Public Debt

We will now show how our theory of consumption can be used to derive the effects of the government's tax and debt policies on private consumption demand. This will give us an opportunity to highlight the importance of the way in which consumers form their expectations about the future. It will also enable us to illustrate the important distinction between temporary and permanent changes in tax policy.

Temporary versus Permanent Tax Cuts

To focus on tax policy, we start by rewriting the consumption function (17) as

$$C_1 = \theta \left(Y_1^L - T_1 + \frac{Y_2^L - T_2}{1+r} + V_1 \right). \tag{22}$$

Suppose now that the government wishes to stimulate consumption demand by cutting net taxes T_1 during period 1, say, because the economy is in recession. For the moment, we assume that consumers expect the tax cut to be *temporary*, perhaps because the government has announced that it may have to raise taxes again once the recession is over. In that case consumers will expect T_2 to be unchanged, and according to (22) the immediate impact of the temporary tax cut will then be given by the derivative

$$\frac{\partial C_1}{\partial T_1} = -\theta. \tag{23}$$

In other words, if the government temporarily cuts taxes by one unit, current consumption will go up by θ units.[11] Thus the temporary tax cut will indeed succeed in stimulating current consumption, but recalling from (16) that $\theta < 1$, we also see that part of the increase in disposable income will be saved for future consumption. Of course, this saving reflects the consumer's desire to smooth consumption over time.

For comparison, suppose instead that consumers expect the tax cut to be *permanent* so that T_2 goes down by the same amount as T_1. From (22) we then find the effect on current consumption to be

$$\frac{\partial C_1}{\partial T_1} + \frac{\partial C_1}{\partial T_2} = -\theta \left(1 + \frac{1}{1+r} \right). \tag{24}$$

Not surprisingly, we see from (23) and (24) that *a permanent tax cut will have a stronger impact on current consumption than a temporary tax cut.* In the benchmark case where

[11] Note that this is a *partial* analysis. If the tax cut succeeds in stimulating economic activity, it will indirectly cause a further rise in consumption by raising aggregate labour income Y_1^L. However, our focus here is on the immediate direct impact on consumption, for a given level of activity.

the real interest rate r equals the rate of time preference ϕ so that the consumer wants to consume equal amounts in the two periods, you may easily verify from (16) that $\theta = (1+r)/(2+r)$. In that case it follows from (24) that

$$\frac{\partial C_1}{\partial T_1} + \frac{\partial C_1}{\partial T_2} = -1 \qquad \text{for} \quad r = \phi.$$

In other words, when the consumer wants to smooth his consumption perfectly over time, he will consume all of his period 1 tax cut during that period. There is no need for him to save any part of the increase in current disposable income to smooth consumption, since his expected future disposable income has gone up by a similar amount as his current net income.

The Government Budget Constraint

In the analysis above we have not specified whether the tax cuts are associated with changes in public spending. As we shall argue below, the effects of a tax cut on private consumption may depend on whether or not it reflects a cut in public expenditure. To see this, we must make a slight detour to study the government's budget constraint. Let us assume for the moment that the government plans for the same time horizon as the household sector (later we will discuss the implications of relaxing this assumption). Like before, we will divide this time horizon into two periods representing the 'present' (period 1) and the 'future' (period 2). At the beginning of period 1 the government starts out with a given real stock of public debt D_1 which is predetermined by the accumulated historical government budget deficits. In period 1 (at the beginning) the government spends a real amount G_1 on public consumption and collects a real net tax revenue T_1, where T_1 measures taxes net of government transfer payments (transfers are simply treated as negative taxes). If public spending exceeds tax revenues, the government must issue new government bonds in the real amount $\Delta D = G_1 - T_1$ in period 1.[12] Since debt carries interest, the government

[12] We abstract from money printing (seigniorage) as a way of financing government consumption, since

will then end up with a total amount of real debt $D_2 = (1+r)(D_1 + \Delta D)$ at the beginning of period 2. Hence the government budget constraint for period 1 may be written as

$$D_2 = (1+r)(D_1 + G_1 - T_1). \tag{25}$$

In period 2 the government must run a budget surplus $T_2 - G_2$ which is sufficient to enable it to repay the public debt accumulated during the previous period. The government budget constraint for period 2 is therefore given by

$$T_2 = D_2 + G_2. \tag{26}$$

Inserting (25) into (26) and dividing through by $1+r$, we obtain *the intertemporal government budget constraint:*

$$D_1 + G_1 + \frac{G_2}{1+r} = T_1 + \frac{T_2}{1+r} \tag{27}$$

Equation (30) says that the present value of current and future tax revenues must cover the present value of current and future government spending plus the initial government debt. As we shall now see, the intertemporal government budget constraint may have profound implications for the effects of tax policy.

Tax Finance versus Debt Finance of Government Spending

It follows directly from (27) that if the government does not reduce its current or planned future spending - that is, if G_1 and G_2 are unchanged - a tax cut $dT_1 < 0$ in the current period must be followed by a future increase $dT_2 > 0$ in taxes so that

$$dT_1 + \frac{dT_2}{1+r} = 0 \qquad \text{or} \qquad dT_2 = -(1+r)\,dT_1. \tag{28}$$

To understand this relationship between current and future taxes, consider (25) and (26) once again. According to (25), if the government cuts current taxes by an amount $|dT_1|$

this plays a very little role in developed countries.

without reducing current public spending G_1, the stock of real public debt will have risen by the amount $dD_2 = (1+r)\,|dT_1|$ at the beginning of period 2. If the government does not cut back on its period 2 consumption G_2, it follows from (26) that it will have to raise taxes in period 2 by the amount $dT_2 = (1+r)\,|dT_1|$ to pay for the principal and interest on the extra debt created by the tax cut in period 1.

Suppose now that consumers have *rational expectations* in the sense that they look forward and understand the implications of the intertemporal government budget constraint. In that case they will realize that if the government reduces current taxes *without* reducing current or planned future public spending, the present value of future taxes will have to increase by as much as the cut in current taxes. It then follows from (22) and (28) that

$$dC_1 = -\theta\left(dT_1 + \frac{dT_2}{1+r}\right) = 0 \tag{29}$$

The implication of (29) is striking: a cut in current taxes which is *not* accompanied by a cut in present or planned future public spending will have *no* effect on private consumption! In other words, a switch from tax finance to debt finance of current public spending leaves private consumption unaffected, since consumers realize that the tax cut today will be offset by the future tax increase needed to service the additional government debt. Indeed, because the present value of their lifetime tax burden is unchanged, consumers know that their net human wealth is also unchanged, and hence they feel unable to afford an increase in current consumption. Instead they will save the entire current tax cut and invest it in the capital market. This increase in current private saving will raise consumer cash receipts in period 2 by the amount $(1+r)\,|dT_1|$ which is just sufficient to pay for the higher future tax bill. Hence future private consumption is also unchanged. The increased supply of government bonds in period 1 is thus matched by a similar increase in private demand for bonds, and the higher period 2 taxes are matched by an increase in private sector income from bond holdings.

The observation that tax finance and debt finance of government spending are in principle equivalent was first made as early as 1820 by the classical British economist David Ricardo. In a treatise on public debt Ricardo discussed three alternative ways of financing a war. For concreteness, he assumed that the war would generate military expenditure of 20 million pounds per year. One financing option would be to impose additional taxes amounting to 20 million pounds per year until the end of the war. Alternatively, the government could borrow 20 million pounds every year during the war and increase tax collections by just 1 million pounds each year to cover the interest payments on a 20 million loan, assuming an interest rate of 5 percent. In this case the public debt would continue to rise until the war was over and would never be repaid, and taxes would be permanently higher. The third possibility considered by Ricardo was one where the war was mainly debt-financed, but where taxes would be raised by 1.2 million pounds per year for every 20 million borrowed, enabling the government to pay off its debt in 45 years. Under the assumptions made, the present value of tax payments would be the same under the three modes of finance. As Ricardo put it: "In point of economy, there is no real difference in either of the modes; for twenty millions in one payment, one million per annum for ever, or 1,200,000 pounds for 45 years, are precisely of the same value".[13]

Because of this statement by Ricardo, the claim that taxes and debt are equivalent methods of public finance is referred to as the *Ricardian Equivalence Theorem*. Before you dismiss the theorem as being utterly unrealistic, take a look at Figure 17.9. In that figure we have plotted private and public saving in Denmark as a percentage of GDP against time. As we have seen above, the Ricardian Equivalence Theorem implies that whenever the public sector reduces its saving, thereby increasing new issues of public debt (or reducing the rate at which public debt is retired), we should observe an offsetting increase in private

[13]See Ricardo's paper on the 'Funding System', p. 186, volume IV of *The Works and Correspondence of David Ricardo,* edited by Piero Sraffa, Cambridge University Press, 1951. The subtle phrase: "In point of economy ..." has been interpreted to mean something like: "From a rational economic perspective...".

saving, and vice versa. Figure 17.9 shows that there is indeed a clear tendency for private and public saving to move in opposite directions, suggesting that Ricardian equivalence may not be that unrealistic after all.

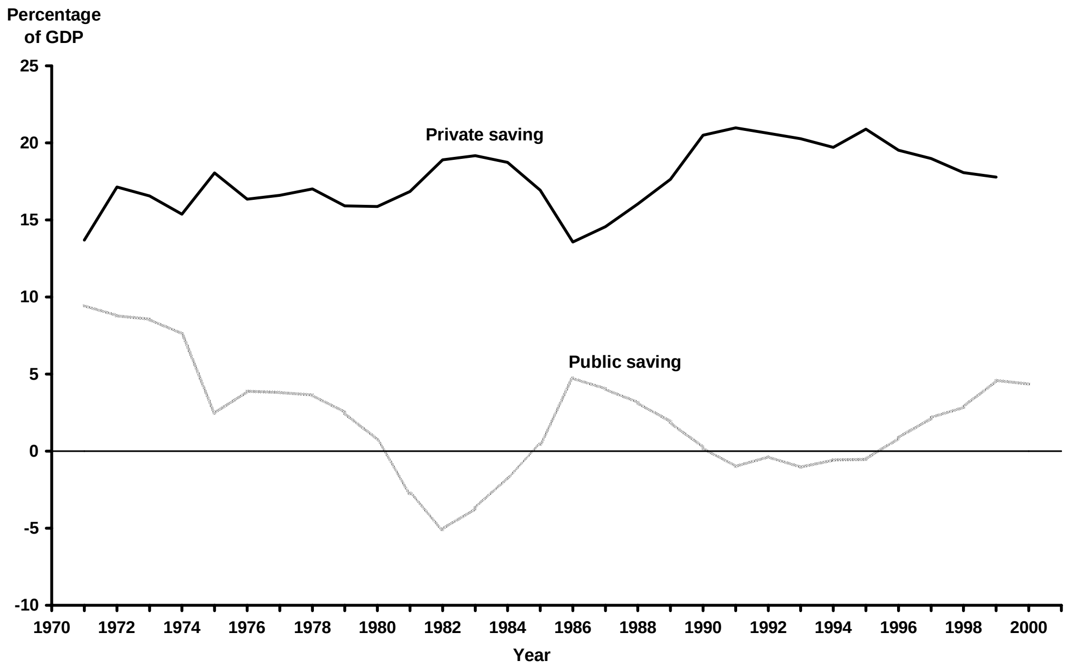

Figure 17.9: Private and public saving in Denmark

Note: Private saving is gross saving of companies and households relative to GDP. Public saving is gross public saving relative to GDP.
Source: ADAM database, Statistics Denmark

However, while Figure 17.9 does not seem inconsistent with Ricardian Equivalence, it does not 'prove' that the theorem is correct, since there may be other explanations for the observed relationship between private and public saving. One potential Keynesian explanation is that Figure 17.9 simply reflects the so-called automatic stabilizers built into the public budget: if the private propensity to save goes up for some reason, the resulting fall in consumption demand will tend to generate a fall in economic activity which in turn will increase the public budget deficit by automatically reducing government revenue from taxes on income and consumption and by automatically increasing public expenses on unemployment benefits. Alternatively, if the economy is hit by some other contractionary shock, say, a drop in investment, the resulting fall in economic activity may increase the private propensity to save at the same time as it increases the budget deficit via the automatic stabilizers. Thus, Figure 17.10 suggests that reductions in the rate of employment are associated with reductions in the average propensity to consume, perhaps because higher unemployment reduces expected future income growth or increases uncertainty about the future.[14]

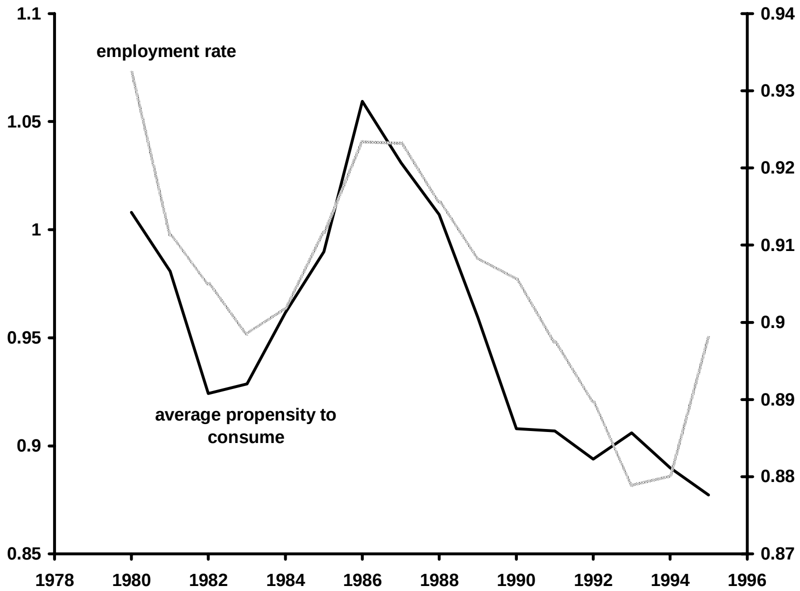

Figure 17.10: The rate of employment and the average propensity to consume in Denmark

Source: SMEC database, Danish Economic Council

Towards a More Realistic Theory of Consumption

Although Figure 17.9 suggests that the Ricardian Equivalence Theorem cannot be dismissed so easily, most economists remain sceptical of the theorem in its strong form. Indeed, Ricardo himself did not believe in the practical relevance of his theorem. Right after having explained that his three alternative methods of war finance would imply the

[14]Of course, the fact that the employment rate and the propensity to consume move together does not necessarily mean that movements in employment *cause* movements in consumption. Causality could also run in the opposite direction, or the two time series could be driven by a common third factor not included in the figure.

same present value of taxes (cf. the quote above), he proceeded to write: "... but the people who pay the taxes never so estimate them, and therefore do not manage their private affairs accordingly. We are too apt to think that the war is burdensome only in proportion to what we are at the moment called to pay for it in taxes, without reflecting on the probable duration of such taxes".[15] What Ricardo essentially says here is that ordinary taxpayers simply do not have the foresight or sophistication to calculate the present value of their expected future taxes. Hence they do not realize that lower taxes today (more government debt) must mean higher taxes tomorrow as long as the time path of public consumption is unchanged. Casual observation suggests that many people may indeed be myopic in this sense. But even if consumers are not irrational or myopic, there may still be several reasons why taxes and debt are not fully equivalent modes of public finance. Let us briefly consider these reasons.

Finite horizons and intergenerational distribution effects. Our analysis above assumed that the private and public sectors have identical planning horizons. In practice, the state will continue to exist beyond the finite lifetime of the individual consumer. Many current taxpayers (especially the elderly) may therefore rationally expect that some of the future taxes needed to service the existing public debt will be levied on future generations and not on those currently alive. In that case a shift from tax finance to debt finance is a way of shifting (part of) the burden of paying for current government spending onto future generations. A debt-financed tax cut will therefore increase the human wealth and consumption of the currently living generations.

This argument against Ricardian Equivalence may sound plausible, but it was met by an ingenious objection from Chicago economist Robert Barro. He pointed out that parents care about their children. If the government tries to shift part of the tax burden from current to future generations through debt finance, parents may use their tax savings to

[15]Ricardo, op.cit., pp. 186-187.

increase their bequests to compensate their children for the higher future tax burden. When parents internalize the welfare of their children who in turn care about the welfare of *their* children and so on, the current generation will effectively behave as if it has an infinite time horizon. Any attempt by the government to redistribute resources across generations will then be neutralized by offsetting private intergenerational transfers, according to Barro.[16] However, this assumes that all parents actually plan to leave bequests to their children. In reality, some parents may not do so, for example if they believe that their children will be much richer than themselves. Moreover, the Barro argument abstracts from population growth. If population is growing, the future tax burden will be spread out across a larger number of taxpayers, and parents will not have to pass on all of their current tax savings to their children to compensate the latter for the higher future taxes needed to service the higher public debt. Still, Barro's analysis is an important reminder that human mortality as such is not a sufficient reason to dismiss Ricardian equivalence.

Intragenerational redistribution. The macro analysis of the preceding section hides the fact that most taxes are redistributive in nature. For the individual taxpayer or the individual family dynasty, a tax cut today may not be matched by an equivalent present-value tax increase tomorrow, even if the government obeys its aggregate intertemporal budget constraint. If a shift in the tax burden over time also involves a shift in the lifetime tax burdens across different individuals or families, the resulting *intragenerational* redistribution across agents belonging to the same generation may have macroeconomic effects. The reason is that different individuals or families have different characteristics so that redistri-

[16] Robert J. Barro (1974): "Are Government Bonds Net Wealth?", *Journal of Political Economy*, vol. 82, pp. 1095-1117. This seminal article is the authoritative modern statement of the Ricardian Equivalence Theorem. Interestingly, even Barro's sophisticated reasoning was anticipated by Ricardo, for our previous quote from Ricardo's text continues: "It would be difficult to convince a man possessed of 20,000 pounds, or any other sum, that a perpetual payment of 50 pound per annum was equally burdensome with a single tax of 1000 pounds. He would have some vague notion that the 50 pounds per annum would be paid by posterity, and would not be paid by him; but if he leaves his fortune to his son, and leaves it charged with this perpetual tax, where is the difference whether he leaves him 20,000 pounds with the tax, or 19,000 pounds without it?"

bution of income does not necessarily leave aggregate consumption unchanged. However, intragenerational redistribution does not mean that a switch from tax finance to debt finance will necessarily raise current consumption. As a counterexample, if a lot of people fear that they will have to bear a disproportionate share of the higher future tax burden implied by an increase in public debt, a debt-financed cut in current taxes may actually *reduce* current private consumption by reducing expected aggregate human wealth.

Distortionary taxes. The Ricardian Equivalence Theorem assumes that taxes take the form of lump sum payments which are unrelated to individual economic behaviour. In practice a person's tax bill is typically linked to his income or consumption. Taxes on income and consumption may discourage labour supply and savings and induce consumers to change their pattern of consumption. In particular, if a debt-financed cut in current income taxes generates expectations of higher future tax rates to service the higher public debt, consumers may want to increase their current labour supply and reduce their future labour supply to take advantage of the fact that marginal tax rates are lower today than they will be tomorrow. By inducing such *intertemporal substitution in labour supply*, the government may be able to stimulate current economic activity through a switch from tax finance to debt finance of current public spending.

Credit constraints. Our derivation of the Ricardian Equivalence Theorem relied on our assumption that capital markets are perfect so that no consumers are credit-constrained. In reality some consumers may be unable to borrow as much as they would like at the going market interest rate. For example, a person may expect that he will earn much more in the future than today, but if his bank does not have the information necessary to estimate his future earnings potential, it may be reluctant to grant him a credit against his expected future labour income if he cannot produce any collateral. Such a credit-constrained consumer will wish to spend all of his current disposable income here and now, since he would prefer to increase his current consumption at the expense of future consumption if he could only obtain more credit. If the government implements a debt-financed tax cut, credit-constrained consumers will therefore increase their current consumption even if they realize that they will face higher taxes in the future. Via the debt-financed cut in present taxes the government is using its access to the capital market to shift the consumption possibilities of credit-constrained consumers from the future to the present, thereby offsetting the imperfections in the private capital market. In this scenario Ricardian Equivalence breaks down, as you are invited to demonstrate in Exercise 1.

Empirical research has found that current consumption tends to react more strongly to changes in current income than one would expect if consumption were governed only by expected lifetime income (of which current annual or quarterly income is usually only a small fraction). This so-called 'excess sensitivity' of current consumption to current income may reflect that many consumers are indeed credit-constrained and will hence wish to consume all of an increase in their current income.

4 The generalized consumption function

Let us end this chapter by summarizing our theory of private consumption. Although we have just argued that the consumption of liquidity-constrained consumers will depend only on current disposable income, it is also realistic to assume that many consumers will not be credit-constrained in any particular year. For example, people with positive net financial assets are unlikely to be liquidity-constrained, since they can always sell off some of their assets if they want to increase their present consumption relative to planned future consumption. For these individuals we have seen that consumption depends on expected future income as well as on current income. In other words, aggregate consumption depends on current disposable labour income, Y_1^d, and on the expected rate of income growth (g) for those consumers who are not credit-constrained. We have also seen that consumption is affected by the real rate of interest, r, and by the market value of initial private wealth, V_1. We may therefore sum up our theory of consumption in the following generalized consumption function $C(\cdot)$:

$$C_1 = C\left(\underset{(+)}{Y_1^d}, \underset{(+)}{g}, \underset{(?)}{r}, \underset{(+)}{V_1}\right). \quad (30)$$

The signs below the variables in (30) indicate the signs of the partial derivatives implied by our theory. Recall that, because of offsetting income and substitution effects, we cannot say for sure whether a rise in the real interest rate will raise or lower consumption. However, once we reckon that the value of financial and human wealth varies negatively with the interest rate, it becomes more likely that a higher real interest rate will cause consumption to fall, in accordance with popular beliefs.

Notice also how expectations feed into consumption. When consumers are optimistic about the future, they will expect a relatively high rate of income growth, g. This will have a direct positive effect on current consumption. A high expected growth rate will also tend to imply high stock prices and high prices of owner-occupied housing. This will further stimulate consumption via a positive impact on private wealth V_1.

In Chapter 19 we shall see how our consumption function (30) may be combined with our investment function from Chapter 16 to give a theory of aggregate demand for goods and services. But before we can set up our model of aggregate demand and supply, we must study the economy's supply side. This is the subject of the next chapter.

5 Summary

1. Private consumption is by far the largest component of the aggregate demand for goods and services. A satisfactory theory of private consumption must explain the stylized facts that the average propensity to consume is a decreasing function of disposable income in microeconomic cross-section data, whereas it is rougly constant in long run macroeconomic

time series data.

2. The properties of the aggregate consumption function may be derived by studying the behaviour of a representative consumer who must allocate his consumption optimally over time, subject to his intertemporal budget constraint. With perfect capital markets, this constraint implies that the present value of lifetime consumption cannot exceed the sum of the consumer's initial financial and human wealth. Human wealth is the present value of current and future disposable labour income.

3. In the consumer's optimum, the marginal rate of substitution between present and future consumption equals the relative price of future consumption, given by one plus the real rate of interest. When disposable income varies over time, the optimizing consumer will want to smooth the time path of consumption relative to the time path of income. There is evidence that such consumption smoothing does indeed take place.

4. Given the assumption of perfect capital markets, the optimal intertemporal allocation of consumption implies that current consumption is proportional to total current wealth (the sum of financial and human wealth). The propensity to consume current wealth depends on the real interest rate, the consumer's rate of time preference, and on his intertemporal elasticity of substitution, defined as the percentage change in the ratio of future to present consumption implied by a one percent change in the consumer's marginal rate of substitution.

5. A rise in the real interest rate will have offsetting income and substitution effects on the propensity to consume current wealth. If the intertemporal substitution elasticity is greater than one, reflecting a great willingness of consumers to substitute future for present consumption, the substitution effect will dominate. A rise in the real interest rate will then reduce the propensity to consume current wealth. The opposite will happen if the intertemporal substitution elasticity is smaller than one. Even if a change in the interest rate does not significantly affect the propensity to consume a given amount of wealth, it

may reduce current consumption by reducing the present value of future labour income, that is, by reducing human wealth, and by reducing the market value of the consumer's shares and housing wealth.

6. For an optimizing consumer, the average propensity to consume current *income* will vary positively with the ratio of current financial wealth to current income. There is strong empirical evidence that such a positive relationship exists. The rough long-run constancy of the average propensity to consume observed in macroeconomic time series data may be explained by the facts that the wealth-income ratio and the real rate of interest tend to be roughly constant over the long run.

7. The negative correlation between income and the average propensity to consume observed in microeconomic cross-section data may be explained by the fact that, in any given period, many consumers will have a relatively low current income relative to their average income over the life cycle. Such consumers will therefore have a high level of current consumption relative to their current income, because they expect higher future incomes, or because they have accumulated wealth by saving out of higher past incomes.

8. A tax cut which is expected to be permanent will have a stronger positive impact on current consumption than a tax cut which is believed to be temporary. When the real interest rate equals the rate of time preference, a permanent tax cut will induce a corresponding rise in current consumption.

9. The government's intertemporal budget constraint implies that the present value of current and future taxes must be sufficient to cover the present value of current and future government spending plus the initial stock of government debt. For given levels of current and future government consumption, a tax cut today must therefore be offset by a future tax increase of equal present value.

10. If consumers have rational expectations they will realize the implications of the intertemporal government budget constraint. This means that a cut in current (lump

sum) taxes which is not accompanied by a cut in present or future public spending will have no effect on private consumption: consumers will save all of the current tax cut to be able to finance the higher future taxes without having to reduce future consumption. This equivalence between tax finance and debt finance of current public spending is referred to as Ricardian Equivalence.

11. In practice, consumers are unlikely to save the full amount of a current tax cut, even if they realize that lower taxes today must imply higher taxes in the future. First of all, consumers may believe that some of the future taxes will be levied on future generations. Second, some consumers may be credit-constrained. A switch from current to future taxes will help these individuals to achieve a desired rise in current consumption at the expense of future consumption. The use of redistributive and distortionary taxes also means that a switch from tax finance to debt finance of current public spending is likely to have real effects on current consumption and labour supply.

12. The theory of private consumption is summarized in the generalized consumption function which states that aggregate consumption is an increasing function of current disposable income, of the expected future growth rate of income, and of the current ratio of financial wealth to income. A rise in the real interest rate has a theoretically ambiguous effect, although it is likely to reduce current consumption due to its negative impact on human and financial wealth.

6 Exercises

Exercise 1. Tax policy and consumption with credit-constrained consumers

In part 2 of this chapter we saw that the Ricardian Equivalence Theorem rests on the assumption of perfect capital markets. In this exercise you are asked to analyze the effects of a switch from tax finance to debt finance of public consumption when some consumers are credit-constrained.

For concreteness, suppose that all consumers in the economy earn the same labour income and pay the same amount of taxes in each period, but that one group of consumers ('the poor') enters the economy with a zero level of initial wealth at the beginning of period 1, whereas the remaining consumers ('the rich') start out with a level of initial wealth equal to V_1. Moreover, suppose that disposable labour income in period 1 is so low that the poor would like to borrow during that period, but that the banks are afraid of lending them money because they cannot provide any collateral. In that case the poor will be credit-constrained during period 1, and the consumption of a poor person during that period, C_1^p, will then be given by the budget constraint

$$C_1^p = Y_1^L - T_1 \tag{1}$$

A rich consumer does not face any borrowing constraint, and his optimal consumption in period 1, C_1^r, will therefore be given by the consumption function (16) from the main text, that is

$$C_1^r = \theta \left(V_1 + Y_1^L - T_1 + \frac{Y_2^L - T_2}{1+r} \right), \qquad 0 < \theta < 1 \tag{2}$$

Suppose that the total population size is equal to 1 (we can always normalize population size in this way by appropriate choice of our units of measurement). Suppose further that a fraction μ of the total population is 'poor' in the sense of having no initial wealth.

1. Derive the economy's aggregate consumption function for period 1, that is, derive an expression for total consumption $C_1 = \mu C_1^p + (1-\mu) C_1^r$. Derive an expression for the economy's marginal propensity to consume current disposable income, $\partial C_1/\partial \left(Y_1^L - T_1\right)$. Compare this expression to the value of the marginal propensity to consume in an economy with no credit-constrained consumers. Explain the difference.

Suppose now that the government enacts a debt-financed reduction of current taxes T_1 by one unit, without cutting current or planned future public consumption. Assume further that the government and the private sector have the same planning horizons, and that the private sector understands that the tax cut today will have to be matched by a tax hike tomorrow of the same present value, due to the intertemporal government budget constraint. In other words, suppose that all consumers realize that

$$dT_1 + \frac{dT_2}{1+r} = 0 \tag{3}$$

2. Derive the effect of the switch from tax finance to debt finance on aggregate private consumption in period 1. Compare to the situation with no credit-constrained consumers and explain the difference. Discuss whether the increase in public debt in period 1 will improve the lifetime welfare of consumers?

Instead of financing the tax cut in period 1 by debt, the government may finance the tax reduction by cutting public consumption. In the question below we will distinguish between two scenarios. In the first scenario the fall in public consumption is expected to be *temporary*, that is, $dG_1 = dT_1$ and $dG_2 = dT_2 = 0$. In the second scenario the cut in public consumption is expected to be *permanent* so that $dG_2 = dG_1 = dT_1 = dT_2$.

3. Derive the effect on current private consumption C_1 of a temporary tax cut financed by a temporary cut in public consumption. Compare this to the effect on C_1 of a permanent tax cut financed by a permanent cut in public consumption (in the latter case you may assume that $r = \phi$). Explain the difference between your expressions. Does it make any difference for the effects of temporary and permanent tax cuts whether consumers are credit-constrained or not?

Exercise The consumption function with endogenous labour supply

In the main text of this chapter we made the simplifying assumption that the consumer's labour income is exogenously given to him, say, because wage rates as well as working hours are regulated by collective bargaining agreements. This exercise invites you to derive the consumption function when the representative consumer may freely choose his preferred number of working hours h, while still taking the real wage rate w as given from the market.

To simplify, we will assume that the consumer is retired from the labour market during the second period of his life, earning labour income only during the first period. In real terms, his budget constraints for the two periods of life may then be written as

$$V_2 = (1+r)\left[V_1 + w(1-\tau)h - C_1\right] \quad (1)$$

$$C_2 = V_2 \quad (2)$$

where τ is a proportional labour income tax so that $w(1-\tau)h$ is disposable labour income. For concreteness, we assume that the consumer's lifetime utility is given by the function

$$U = \ln C_1 + \beta \ln(1-h) + \frac{\ln C_2}{1+\phi} \quad (3)$$

where ϕ is the rate of time preference. In (3) we have assumed that the total time available to the consumer in the first period is equal to 1. Hence the magnitude $1-h$ is the amount of leisure enjoyed during period 1, and β is a parameter indicating the strength of the consumer's preference for leisure. Furthermore, the after-tax real wage rate $w(1-\tau)$ may be called the consumer's *potential* labour income (or the market value of his time endowment), since it measures the amount of wage income he could earn if he worked all the time. Note that $w(1-\tau)$ may also be seen as the 'price' of leisure, since it measures the net income foregone by the consumer if he chooses to consume one more unit of leisure.

1. Derive and interpret the consumer's intertemporal budget constraint. Does lifetime consumption depend on actual or on potential labour income? (hint: for later purposes it may be useful for you to rewrite the lifetime budget constraint so that the value of the consumption of leisure, $w(1-\tau)(1-h)$, appears on the left-hand side as part of total lifetime consumption).

2. Derive an expression for the consumer's optimal consumption during period 1, assuming that he wishes to maximize his lifetime utility (3) (hint: rewrite the utility function and the intertemporal budget constraint by introducing the variable $F \equiv 1-h$, where F is leisure, and use the intertemporal budget constraint to eliminate C_2 from the utility function (3). Then derive the first-order conditions for the optimal choice of C_1 and F and use these expressions to derive the expression for the optimal value of C_1). Compare your consumption function to the consumption function (16) derived in the the main text of the chapter and comment on the similarities and differences.

3. Derive an expression for the consumer's optimal labour supply h during period 1 (recall that $h = 1-F$). How does labour supply depend on the labour income tax rate τ? Give an intuitive explanation.

4. Suppose now that the consumer's desired working hours h exceed the amount of hours $\overline{h}$ which he is actually allowed to work as a member of his trade union. In that case he will work the maximum hours allowed, and his disposable labour income in period 1 will be $w(1-\tau)\overline{h}$. Derive his optimal period 1 consumption level in this case and compare to the consumption function derived in Question 2. Comment on the difference.

Exercise 3. Wage taxes versus consumption taxes versus wealth taxes

In the main text of the chapter we assumed for simplicity that taxes took the form of lump sum payments which were unrelated to the consumer's behaviour. Now we assume instead that the consumer must pay a proportional tax τ^w on his wage income, a proportional value-added tax τ^c on all his consumption expenditure, and potentially also a one-time tax τ^v on his initial wealth. This exercise asks you to consider some similarities and differences between these taxes.

We assume that the consumer's working hours are exogenously given to him and equal to 1 so that his total pre-tax labour income in period 1 is equal to the real wage rate w. We also assume that he is retired from the labour market during period 2. We start out in an initial situation without any taxes on wealth. The consumer budget constraints for the two periods of life are then given by

$$V_2 = (1+r)\left[V_1 + (1-\tau^w)\,w - (1+\tau^c)\,C_1\right], \qquad 0 < \tau^w < 1, \qquad \tau^c \geq 0 \tag{1}$$

$$(1+\tau^c)\,C_2 = V_2 \tag{2}$$

1. Derive the consumer's intertemporal budget constraint and comment on your expression. From the consumer's perspective, what is the similarity and the difference between the wage tax and the consumption tax?

The consumer has the lifetime utility function

$$U = \ln C_1 + \frac{\ln C_2}{1+\phi}, \qquad \phi > 0 \tag{3}$$

2. Derive the consumer's optimal values of C_1 and C_2 and comment on your expressions.

Suppose now that the government needs to raise more tax revenue to finance additional public spending. The government presents two alternative proposals in parliament. The first proposal implies that the consumption tax rate (the VAT) will be raised from 20% to 25% whereas the wage tax rate will be kept unchanged at 50%. The second proposal implies that the consumption tax rate is maintained at 20%, whereas the wage tax rate is raised from 50% to 52%. In addition, the second proposal includes a one-time proportional tax of 4 percent on existing initial wealth V_1. The government stresses that this is a once-and-for-all wealth tax which will not be imposed on future wealth V_2.

3. Would consumers prefer one of the government's proposals to the other one? What difference does it make to your answer whether the wealth tax is a one-time levy or a permanent tax?