Rational Expectations, Consumption and Asset Pricing

Elementary Keynesian macro theory assumes that households make consumption decisions based only on their current disposable income. In reality, of course, people have to base their spending decisions not just on today's income but also on the money they expect to earn in the future. During the 1950s, important research by Ando and Modigliani (the Life-Cycle Hypothesis) and Milton Friedman (the Permanent Income Hypothesis) presented significant evidence that people plan their expenditures in system pattern, smoothing consumption over time even when their incomes fluctuated.

In these notes, we will use the techniques developed in the last topic to derive a rational expectations version of the Permanent Income Hypothesis. We will use this model to illustrate some pitfalls in using econometrics to assess the effects of policy changes. We will discuss empirical tests of this model and present some more advanced topics. In particular, we will discuss the link between consumption spending and the return on various financial assets.

The Household Budget Constraint

We start with an identity describing the evolution of the stock of assets owned by households. Letting A_t be household assets, Y_t be labour income, and C_t stand for consumption spending, this identity is

$$A_{t+1} = (1 + r_{t+1}) (A_t + Y_t - C_t) \tag{1}$$

where r_{t+1} is the return on household assets at time $t + 1$. Note that Y_t is *labour* income (income earned from working) not total income because total income also includes the capital income earned on assets (i.e. total income is $Y_t + r_{t+1}A_t$.) Note, we are assuming that Y_t is take-home labour income, so it can considered net of taxes.

As with the equation for the return on stocks, this can be written as a first-order difference equation in our standard form

$$A_t = C_t - Y_t + \frac{A_{t+1}}{1 + r_{t+1}} \tag{2}$$

We will assume that agents have rational expectations. Also, in this case, we will assume that the return on assets equals a constant, r. This implies

$$A_t = C_t - Y_t + \frac{1}{1+r} E_t A_{t+1} \tag{3}$$

Using the same repeated substitution methods as before this can be solved to give

$$A_t = \sum_{k=0}^{\infty} \frac{E_t \left(C_{t+k} - Y_{t+k}\right)}{(1+r)^k} \tag{4}$$

Note that we have again imposed the condition that the final term in our repeated substitution $\frac{E_t A_{t+k}}{(1+r)^k}$ goes to zero as k gets large. Effectively, this means that we are assuming that people consume some of their capital income (i.e. that assets are used to finance a level of consumption C_t that is generally larger than labour income Y_t). If this is the case, then this term tends to zero.

One way to understand this equation comes from re-writing it as

$$\sum_{k=0}^{\infty} \frac{E_t C_{t+k}}{(1+r)^k} = A_t + \sum_{k=0}^{\infty} \frac{E_t Y_{t+k}}{(1+r)^k} \tag{5}$$

This is usually called the *intertemporal budget constraint*. It states that the present value sum of current and future household consumption must equal the current stock of financial assets plus the present value sum of current and future labour income.

A consumption function relationship can be derived from this equation by positing some theoretical relationship between the expected future consumption values, $E_t C_{t+k}$, and the current value of consumption. This is done by appealing to the optimising behaviour of the consumer.

Optimising Behaviour by the Consumer

We will assume that consumers wish to maximize a welfare function of the form

$$W = \sum_{k=0}^{\infty} \left(\frac{1}{1+\beta} \right)^k U\left(C_{t+k}\right) \tag{6}$$

where $U\left(C_t\right)$ is the instantaneous utility obtained at time t, and β is a positive number that describes the fact that households prefer a unit of consumption today to a unit tomorrow. If the future path of labour income is known, consumers who want to maximize this welfare function subject to the constraints imposed by the intertemporal budget constraint must solve the following Lagrangian problem:

$$L\left(C_{t,}C_{t+1}, ...\right) = \sum_{k=0}^{\infty} \left(\frac{1}{1+\beta} \right)^k U\left(C_{t+k}\right) + \lambda \left[A_t + \sum_{k=0}^{\infty} \frac{Y_{t+k}}{\left(1+r\right)^k} - \sum_{k=0}^{\infty} \frac{C_{t+k}}{\left(1+r\right)^k} \right] \tag{7}$$

For every current and future value of consumption, C_{t+k}, this yields a first-order condition of the form

$$\left(\frac{1}{1+\beta} \right)^k U'\left(C_{t+k}\right) - \frac{\lambda}{\left(1+r\right)^k} = 0 \tag{8}$$

For $k = 0$, this implies

$$U'\left(C_t\right) = \lambda \tag{9}$$

For $k = 1$, it implies

$$U'\left(C_{t+1}\right) = \left(\frac{1+\beta}{1+r} \right) \lambda \tag{10}$$

Putting these two equations together, we get the following relationship between consumption today and consumption tomorrow:

$$U'\left(C_t\right) = \left(\frac{1+r}{1+\beta} \right) U'\left(C_{t+1}\right) \tag{11}$$

When there is uncertainty about future labour income, this optimality condition can just be re-written as

$$U'\left(C_t\right) = \left(\frac{1+r}{1+\beta} \right) E_t\left[U'\left(C_{t+1}\right)\right] \tag{12}$$

This implication of the first-order conditions for consumption is sometimes known as an *Euler equation*.

In an important 1978 paper, Robert Hall proposed a specific case of this equation.[1] Hall's special case assumed that

$$U(C_t) = aC_t + bC_t^2 \tag{13}$$

$$r = \beta \tag{14}$$

In other words, Hall assumed that the utility function was quadratic and that the real interest rate equalled the household discount rate. In this case, the Euler equation becomes

$$a + 2bC_t = E_t\left[a + 2bC_{t+1}\right] \tag{15}$$

which simplifies to

$$C_t = E_t C_{t+1} \tag{16}$$

This states that the optimal solution involves next period's expected value of consumption equalling the current value. Because, the Euler equation holds for all time periods, we have

$$E_t C_{t+k} = E_t C_{t+k+1} \qquad k = 1, 2, 3, ... \tag{17}$$

So, we can apply repeated iteration to get

$$C_t = E_t\left(C_{t+k}\right) \qquad k = 1, 2, 3, ... \tag{18}$$

In other words, all future expected values of consumption equal the current value. Because it implies that changes in consumption are unpredictable, this is sometimes called the *random walk* theory of consumption.

[1] "Stochastic Implications of the Life-Cycle Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, December 1978.

The Rational Expectations Permanent Income Hypothesis

Hall's random walk hypothesis has attracted a lot of attention in its own right, but rather than focus on what should be unpredictable (changes in consumption), we are interested in deriving an explicit formula for what consumption should equal.

To do this, insert $E_t C_{t+k} = C_t$ into the intertemporal budget constraint, (5), to get

$$\sum_{k=0}^{\infty} \frac{C_t}{(1+r)^k} = A_t + \sum_{k=0}^{\infty} \frac{E_t Y_{t+k}}{(1+r)^k} \tag{19}$$

Now we can use the geometric sum formula to turn this into a more intuitive formulation:

$$\sum_{k=0}^{\infty} \frac{1}{(1+r)^k} = \frac{1}{1-\frac{1}{1+r}} = \frac{1+r}{r} \tag{20}$$

So, Hall's assumptions imply the following equation, which we will term the *Rational Expectations Permanent Income Hypothesis*:

$$C_t = \frac{r}{1+r} A_t + \frac{r}{1+r} \sum_{k=0}^{\infty} \frac{E_t Y_{t+k}}{(1+r)^k} \tag{21}$$

This equation is a rational expectations version of the well-known *permanent income hypothesis* (I will use the term RE-PIH below) which states that consumption today depends on a person's expected lifetime sequence of income.

Let's look at this equation closely. It states that the current value of consumption is driven by three factors:

- The expected present discounted sum of current and future labour income.
- The current value of household assets. This "wealth effect" is likely to be an important channel through which financial markets affect the macroeconomy.

- The expected return on assets: This determines the coefficient, $\frac{r}{1+r}$, that multiplies both assets and the expected present value of labour income. In this model, an increase in this expected return raises this coefficient, and thus boosts consumption.

A Concrete Example: Constant Expected Growth in Labour Income

This RE-PIH model can be made more concrete by making specific assumptions about expectations concerning future growth in labour income. Suppose, for instance, that households expect labour income to grow at a constant rate g in the future:

$$E_t Y_{t+k} = (1+g)^k Y_t \tag{22}$$

This implies

$$C_t = \frac{r}{1+r} A_t + \frac{r Y_t}{1+r} \sum_{k=0}^{\infty} \left(\frac{1+g}{1+r} \right)^k \tag{23}$$

As long as $g < r$ (and we will assume it is) then we can use the geometric sum formula to simplify this expression

$$\sum_{k=0}^{\infty} \left(\frac{1+g}{1+r} \right)^k = \frac{1}{1 - \frac{1+g}{1+r}} \tag{24}$$

$$= \frac{1+r}{r-g} \tag{25}$$

This implies a consumption function of the form

$$C_t = \frac{r}{1+r} A_t + \frac{r}{r-g} Y_t \tag{26}$$

Note that the higher is expected future growth in labour income g, the larger is the coefficient on today's labour income and thus the higher is consumption.

The Lucas Critique

The fact that the coefficients of so-called *reduced-form* relationships, such as the consumption function equation (26), depend on expectations about the future is an important theme in modern macroeconomics. In particular, in a famous 1976 paper, rational expectations pioneer Robert Lucas pointed out that the assumption of rational expectations implied that these coefficients would change if expectations about the future changed.[2] In our example, the MPC from current income will change if expectations about future growth in labour income change.

Lucas's paper focused on potential problems in using econometrically-estimated reduced-form regressions to assess the impact of policy changes. He pointed out that changes in policy may change expectations about future values of important variables, and that these changes in expectations may change the coefficients of reduced-form relationships. This type of problem can limit the usefulness for policy analysis of reduced-form econometric models based on historical data. This problem is now known as the *Lucas critique* of econometric models.

To give a specific example, suppose the government is thinking of introducing a temporary tax cut on labour income. As noted above, we can consider Y_t to be after-tax labour income, so it would be temporarily boosted by the tax cut. Now suppose the policy-maker wants an estimate of the likely effect on consumption of the tax cut. They may get their economic advisers to run a regression of consumption on assets and after-tax labour income. If, in the past, consumers had generally expected income growth of g, then the econometric regressions will report a coefficient of approximately $\frac{r}{r-g}$ on labour income. So, the economic adviser

[2] Robert Lucas, "Econometric Policy Evaluation: A Critique," *Carnegie-Rochester Series on Public Policy*, Vol. 1, pages 19-46, 1976.

might conclude that for each extra dollar of labour income produced by the tax cut, there will be an increase in consumption of $\frac{r}{r-g}$ dollars.

However, if households have rational expectations and operate according to equation (21) then the true effect of the tax cut could be a lot smaller. For instance, if the tax cut is only expected to boost this period's income, and to disappear tomorrow, then each dollar of tax cut will produce only $\frac{r}{1+r}$ dollars of extra consumption. The difference between the true effect and the economic advisor's supposedly "scientific" regression-based forecast could be substantial. For instance, plugging in some numbers, suppose $r = 0.06$ and $g = 0.02$. In this case, the economic advisor concludes that the effect of a dollar of tax cuts is an extra 1.5 $(=\frac{.06}{.06-.02})$ dollars of consumption. In reality, the tax cut will produce only an extra 0.057 $(=\frac{.06}{1.06})$ dollars of extra consumption. This is a big difference.

The Lucas critique has played an important role in the increased popularity of rational expectations economics. Examples like this one show the benefit in using a formulation such as equation (21) that explicitly takes expectations into account, instead of relying only on reduced-form econometric regressions.

Implications for Fiscal Policy: Ricardian Equivalence

Like households, governments also have budget constraints. Here we consider the implications of these constraints for consumption spending in the Rational Expectations Permanent Income Hypothesis. First, let us re-formulate the household budget constraint to explicitly incorporate taxes. Specifically, let's write the period-by-period constraint as

$$A_{t+1} = (1+r)(A_t + Y_t - T_t - C_t) \tag{27}$$

where T_t is the total amount of taxes paid by households. Taking the same steps as before, we can re-write the intertemporal budget constraint as

$$\sum_{k=0}^{\infty} \frac{E_t C_{t+k}}{(1+r)^k} = A_t + \sum_{k=0}^{\infty} \frac{E_t \left(Y_{t+k} - T_{t+k}\right)}{(1+r)^k} \tag{28}$$

Now let's think about the government's budget constraint. The stock of public debt, D_t evolves over time according to

$$D_{t+1} = (1+r)\left(D_t + G_t - T_t\right) \tag{29}$$

where G_t is government spending and T_t is tax revenue. Applying the repeated-substitution method we can obtain an intertemporal version of the government's budget constraint.

$$\sum_{k=0}^{\infty} \frac{E_t T_{t+k}}{(1+r)^k} = D_t + \sum_{k=0}^{\infty} \frac{E_t G_{t+k}}{(1+r)^k} \tag{30}$$

This states that the present discounted value of tax revenue must equal the current level of debt plus the present discounted value of government spending. In other words, in the long-run, the government must raise enough tax revenue to pay off its current debts as well as its current and future spending.

Consider the implications of this result for household decisions. If households have rational expectations, then they will understand that the government's intertemporal budget constraint, equation (30), pins down the present value of tax revenue. In this case, we can substitute the right-hand-side of (30) into the household budget constraint to replace the present value of tax revenue. Doing this, the household budget constraint becomes

$$\sum_{k=0}^{\infty} \frac{E_t C_{t+k}}{(1+r)^k} = A_t - D_t + \sum_{k=0}^{\infty} \frac{E_t \left(Y_{t+k} - G_{t+k}\right)}{(1+r)^k} \tag{31}$$

Consider now the implications of this result for the impact of a temporary cut in taxes. Before, we had discussed how a temporary cut in taxes should have a small effect. This equation gives us an even more extreme result — unless governments plan to change the profile of government spending, then a cut to taxes today has no impact at all on consumption spending. This is because households anticipate that lower taxes today will just trigger higher taxes tomorrow.

This result – that rational expectations implied that a deficit-financed cut in taxes should have no impact on consumption – was first presented by Robert Barro in a famous 1974 paper.[3] It was later pointed out that some form of this result was alluded to in David Ricardo's writings in the nineteenth century. Economists love fancy names for things, so the result is now often referred to as *Ricardian equivalence*.

Evidence on the RE-PIH

There have been lots of macroeconomic studies on how well the RE-PIH fits the data. One problem worth noting is that there are some important measurement issues when attempting to test the theory. In particular, the model's assumption that consumption expenditures only yield a positive utility flow in the period in which the money is spent clearly does not apply to durable goods, such as cars or computers, which yield a steady flow of utility. For this reason, most empirical research has focused only on spending on nondurables (e.g. food) and services, with a separate literature focusing on spending on consumer durables.

There are various reasons why the RE-PIH may not hold. Firstly, it assumes that it is always feasible for households to "smooth" consumption in the manner predicted by the theory. For example, even if you anticipate earning lots of money in the future and would like to have a high level of consumption now, you may not be able to find a bank to fund a lavish lifestyle now based on your promises of future millions. These kinds of "liquidity

[3]Robert Barro (1974). "Are Government Bonds Net Wealth?" *Journal of Political Economy*, Volume 82(6).

constraints" may make consumption spending more sensitive to their current incomes than the RE-PIH predicts. Secondly, people may not have rational expectations and may not plan their spending decisions in the calculating optimising fashion assumed by the theory.

Following Hall's 1978 paper, the 1980s saw a large amount of research on whether the RE-PIH fitted the data. The most common conclusion was that consumption was "excessively sensitive" to disposable income. In particular, changes in consumption appear to be more forecastable than they should be if Hall's random walk idea was correct. Campbell and Mankiw (1990) is a well-known paper that provides a pretty good summary of these conclusions.[4] They present a model in which a fraction of the households behave according to the RE-PIH while the rest simply consume all of their current income. They estimate the fraction of non-PIH consumers to be about a half. A common interpretation of this result is that liquidity constraints have an important impact on aggregate consumption. (A byproduct of this conclusion would be that financial sector reforms that boost access to credit could have an important impact on consumption spending.)

Evidence on Ricardian Equivalence

There is also a large literature devoted to testing the Ricardian equivalence hypothesis. In addition to the various reasons the RE-PIH itself may fail, there are various other reasons why Ricardian equivalence may not hold. Some are technical points. People don't actually live forever (as we had assumed in the model) and so they may not worry about future tax increases that could occur after they have passed away; taxes take a more complicated form than the simple lump-sum payments presented above; the interest rate in the government's

[4] John Campbell and Gregory Mankiw (1990). "Permanent Income, Current Income, and Consumption," *Journal of Business and Economic Statistics*

budget constraint may not be the same as the interest rate in the household's constraint. (You can probably think of a few more.) More substantively, people may often be unable to perceive whether tax changes are temporary or permanent. Most of the macro studies on this topic (in particular those that use Vector Autoregressions) tend to find the effects of fiscal policy are quite different from the Ricardian equivalence predictions. Tax cuts and increases in government spending tend to boost the economy.

Perhaps the most interesting research on this area has been the use of micro data to examine the effect of changes in taxes that are explicitly predictable and temporary. One recent example is the paper by Parker, Souleles, Johnson and Robert McClelland which examines the effect of tax rebates provided to U.S. taxpayers in 2008.[5] This programme saw the U.S. government send once-off payments to consumers in an attempt to stimulate the economy. Since these payments were being financed by expanding the government deficit, Ricardian equivalence predicts that consumers should not have responded. Parker et al, however, found the opposite using data from the Consumer Expenditure Survey. A quick summary:

> We find that, on average, households spent about 12-30% (depending on the specification) of their stimulus payments on nondurable expenditures during the three-month period in which the payments were received. Further, there was also a substantial and significant increase in spending on durable goods, in particular vehicles, bringing the average total spending response to about 50-90% of the payments.

You might suspect that these results are driven largely by liquidity constraints but the

[5] "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review*, 103(6), October 2013.

various microeconomic studies that have examined temporary fiscal policy changes have not always been consistent with this idea. For example, research by Parker (1999) showed the even relatively high-income consumers seemed to spend more in response to transitory changes in their social security taxes (which stop at a certain point in the year when workers reach a maximum threshold point) while Souleles (1999) found "excess sensitivity" results for consumer spending after people received tax rebate cheques.[6] These results show excess sensitivity even among groups of consumers that are unlikely to be liquidity constrained.

At the same time, this doesn't mean that households go on a splurge every time they get a large payment. For example, Hsieh (2003) examines how people in Alaska responded to large anticipated annual payments that they received from a state fund that depends largely on oil revenues. Unlike the evidence on temporary tax cuts, Hsieh finds that Alaskan households respond to these payments in line with the predictions of the Permananet Income Hypothesis, smoothing out their consumption over the year. One possible explanation is that these large and predictable payments are easier for people to understand and plan around and the consequences of spending them too quickly more serious than smaller once-off federal tax changes. There is clearly room for more research in this important area.

[6] Jonathan Parker. "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes," *American Economic Review*, Vol 89 No 4, September 1999. Nicholes Souleles. " The Response of Household Consumption to Income Tax Refunds," *American Economic Review*, Vol 89 No 4, September 1999

Precautionary Savings

I want to return to a subtle point that was skipped over earlier. If we keep the assumption $r = \beta$, then the consumption Euler equation is

$$U'(C_t) = E_t \left[U'(C_{t+1}) \right] \tag{32}$$

You might think that this equation is enough to deliver the property of constant expected consumption. We generally assume declining marginal utility, so function U' is monotonically decreasing. In this case, surely the expectation of next period's marginal utility being the same as this period's is the same as next period's expected consumption level being the same as this period's.

The problem with this thinking is the E_t here is a mathematical expectation, i.e. a weighted average over a set of possible outcomes. And for most functions F generally $E(F(X)) \neq F(E(X))$. In particular, for concave functions—functions like utility functions which have negative second derivatives—a famous result known as Jensen's inequality states that $E(F(X)) < F(E(X))$. This underlies the mathematical formulation of why people are averse to risk: The average utility expected from an uncertain level of consumption is less than from the "sure thing" associated with obtaining the average level of consumption. The sign of the Jensen's inequality result is reversed for concave functions, i.e. those with positive second derivatives.

In this example, we are looking at the properties of $E_t \left[U'(C_{t+1}) \right]$. Whether or not marginal utility is concave or convex depends on its second derivative, so it depends upon the third derivative of the utility function U'''. Most standard utility functions have positive third derivatives implying convex marginal utility and thus $E_t \left[U'(C_{t+1}) \right] > U'(E_t C_{t+1})$. What we can see now is why the quadratic utility function was such a special case. Because this function has $U''' = 0$, its marginal utility is neither concave or convex and the Jensen relationship is an equality. So, in this very particular case, the utility function displays *certainty equivalence*: The uncertain outcome is treated the same way is if people were certain of achieving the average value of consumption.

Here's a specific example of when certainty equivalence doesn't hold.[7] Suppose consumers have a utility function of the form

$$U(C_t) = -\frac{1}{\alpha}\exp\left(-\alpha C_t\right) \tag{33}$$

where exp is the exponential function. This implies marginal utility of the form

$$U'\left(C_t\right) = \exp\left(-\alpha C_t\right) \tag{34}$$

In this case, the Euler equation becomes

$$\exp\left(-\alpha C_t\right) = E_t\left(\exp\left(-\alpha C_{t+1}\right)\right) \tag{35}$$

Now suppose the uncertainty about C_{t+1} is such that it is perceived to have a normal distribution with mean $E_t(C_{t+1})$ and variance σ^2. A useful result from statistics is that if a variable X is normally distributed has mean μ and variance σ^2:

$$X \sim N\left(\mu, \sigma^2\right) \tag{36}$$

then it can be shown that

$$E\left(\exp(X)\right) = \exp\left(\mu + \frac{\sigma^2}{2}\right) \tag{37}$$

[7] This particular example was first presented by Ricardo Caballero (1990), "Consumption Puzzles and Precautionary Savings" *Journal of Monetary Economics*, Volume 25, pages 113-136.

In our case, this result implies that

$$E_t\left(\exp\left(-\alpha C_{t+1}\right)\right) = \exp\left(E_t\left(-\alpha C_{t+1}\right) + \frac{Var\left(-\alpha C_{t+1}\right)}{2}\right) \tag{38}$$

$$= \exp\left(-\alpha E_t\left(C_{t+1}\right) + \frac{\alpha^2\sigma^2}{2}\right) \tag{39}$$

So, the Euler equation can be written as

$$\exp\left(-\alpha C_t\right) = \exp\left(-\alpha E_t\left(C_{t+1}\right) + \frac{\alpha^2\sigma^2}{2}\right) \tag{40}$$

Taking logs of both sides this becomes

$$-\alpha C_t = -\alpha E_t\left(C_{t+1}\right) + \frac{\alpha^2\sigma^2}{2} \tag{41}$$

which simplifies to

$$E_t\left(C_{t+1}\right) = C_t + \frac{\alpha\sigma^2}{2} \tag{42}$$

Even though expected marginal utility is flat, consumption tomorrow is expected to be higher than consumption today. Thus, uncertainty induces an "upward tilt" to the consumption profile. And this upward tilt has an affect on today's consumption: We cannot sustain higher consumption tomorrow without having lower consumption today.

Indeed, it turns out that this result allows us to calculate exactly what the effect of uncertainty is on consumption today. The Euler equation implies that

$$E_t\left(C_{t+k}\right) = C_t + \frac{k\alpha\sigma^2}{2} \tag{43}$$

Inserting this into the intertemporal budget constraint, we get

$$\sum_{k=0}^{\infty} \frac{C_t}{(1+r)^k} + \frac{\alpha\sigma^2}{2} \sum_{k=1}^{\infty} \frac{k}{(1+r)^k} = A_t + \sum_{k=0}^{\infty} \frac{E_t Y_{t+k}}{(1+r)^k} \tag{44}$$

It can be shown (mainly by repeatedly using the well-known geometric sum formula) that

$$\sum_{k=1}^{\infty} \frac{k}{(1+r)^k} = \frac{1+r}{r^2} \tag{45}$$

So, the intertemporal budget constraint simplifies to

$$\sum_{k=0}^{\infty} \frac{C_t}{(1+r)^k} + \frac{1+r}{r^2}\frac{\alpha\sigma^2}{2} = A_t + \sum_{k=0}^{\infty} \frac{E_t Y_{t+k}}{(1+r)^k} \tag{46}$$

and taking the same steps as before, consumption today is

$$C_t = \frac{r}{1+r} A_t + \frac{r}{1+r} \sum_{k=0}^{\infty} \frac{E_t Y_{t+k}}{(1+r)^k} - \frac{\alpha\sigma^2}{2r} \tag{47}$$

This is exactly as before apart from an additional "precautionary savings" term $-\frac{\alpha\sigma^2}{2r}$. The more uncertainty there is, the more lower the current level of consumption will be.

This particular result obviously relies on very specific assumptions about the form of the utility function and the distribution of uncertain outcomes. However, since almost all utility function feature positive third derivatives, the key property underlying the precautionary savings result—marginal utility averaged over the uncertain outcomes being higher than at the average level of consumption—-will generally hold. It is an important result because some of the more important changes in the savings rate observed over time appear consistent with this type of precautionary savings behaviour. So, for example, during the global financial crisis, when there was so much uncertainty about how long the recession would last and what impact it would have, it is very likely that this greater uncertainty depressed consumption.

Incorporating Time-Varying Asset Returns

One simplification that we have made up to now is that consumers expect a constant return on assets. Here, we allow expected asset returns to vary. The first thing to note here is that one can still obtain an intertemporal budget constraint via the repeated substitution method.

This now takes the form

$$\sum_{k=0}^{\infty} \frac{E_t C_{t+k}}{\left(\prod_{m=1}^{k+1} (1 + r_{t+m})\right)} = A_t + \sum_{k=0}^{\infty} \frac{E_t Y_{t+k}}{\left(\prod_{m=1}^{k+1} (1 + r_{t+m})\right)} \tag{48}$$

where $\prod_{n=1}^{h} x_i$ means the product of x_1, x_2 .... x_h. The steps to derive this are identical to the steps used to derive equation (71) in the previous set of notes ("Rational Expectations and Asset Prices").

The optimisation problem of the consumer does not change much. This problem now has the Lagrangian

$$L\left(C_t, C_{t+1}, ...\right) = \sum_{k=0}^{\infty} \left(\frac{1}{1+\beta}\right)^k U\left(C_{t+k}\right) + \lambda \left[A_t + \sum_{k=0}^{\infty} \frac{E_t Y_{t+k}}{\left(\prod_{m=1}^{k+1} (1 + r_{t+m})\right)} - \sum_{k=0}^{\infty} \frac{E_t C_{t+k}}{\left(\prod_{m=1}^{k+1} (1 + r_{t+m})\right)}\right]$$

And instead of the simple Euler equation (12), we get

$$U'\left(C_t\right) = E_t \left[\left(\frac{1 + r_{t+1}}{1+\beta}\right) U'\left(C_{t+1}\right)\right] \tag{49}$$

or, letting

$$R_t = 1 + r_t \tag{50}$$

we can re-write this as

$$U'\left(C_t\right) = E_t \left[\left(\frac{R_{t+1}}{1+\beta}\right) U'\left(C_{t+1}\right)\right] \tag{51}$$

Consumption and Rates of Return on Assets

Previously, we had used an equation like this to derive the behaviour of consumption, given an assumption about the determination of asset returns. However, Euler equations have taken on a double role in modern economics because they are also used to consider the determination of asset returns, taking the path of consumption as given. The Euler equation also takes on greater importance than it might seem based on our relatively simple calculations because, once one extends the model to allow the consumer to allocate their wealth across multiple asset types, it turns out that equation (51) must hold for *all* of these assets. This means that for a set of different asset returns $R_{i,t}$, we must have

$$U'(C_t) = E_t\left[\left(\frac{R_{i,t+1}}{1+\beta}\right)U'(C_{t+1})\right] \tag{52}$$

for each of the assets.

So, for example, consider a pure risk-free asset that pays a guaranteed rate of return next period. The nearest example in the real-world is a short-term US treasury bill. Because there is no uncertainty about this rate of return, call it $R_{f,t}$, these terms can be taken outside the expectation term, and the Euler equation becomes

$$U'(C_t) = \frac{R_{f,t+1}}{1+\beta}E_t\left[U'(C_{t+1})\right] \tag{53}$$

So, the risk-free interest rate should be determined as

$$R_{f,t+1} = \frac{(1+\beta)\,U'(C_t)}{E_t\left[U'(C_{t+1})\right]} \tag{54}$$

To think about the relationship between risk-free rates and returns on other assets, it is useful to use a well-known result from statistical theory, namely

$$E(XY) = E(X)E(Y) + Cov(X,Y) \tag{55}$$

The expectation of a product of two variables equals the product of the expectations plus the covariance between the two variables. This allows one to re-write (52) as

$$U'(C_t) = \frac{1}{1+\beta}\left[E_t\left(R_{i,t+1}\right)E_t\left(U'(C_{t+1})\right) + Cov\left(R_{i,t+1}, U'(C_{t+1})\right)\right] \tag{56}$$

This can be re-arranged to give

$$\frac{(1+\beta)\,U'(C_t)}{E_t\left[U'(C_{t+1})\right]} = E_t\left(R_{i,t+1}\right) + \frac{Cov\left(R_{i,t+1}, U'(C_{t+1})\right)}{E_t\left[U'(C_{t+1})\right]} \tag{57}$$

Note now that, by equation (62), the left-hand-side of this equation equals the risk-free rate. So, we have

$$E_t\left(R_{i,t+1}\right) = R_{f,t+1} - \frac{Cov\left(R_{i,t+1}, U'(C_{t+1})\right)}{E_t\left[U'(C_{t+1})\right]} \tag{58}$$

This equation tells us that expected rate of return on risky assets equals the risk-free rate *minus* a term that depends on the covariance of the risky return with the marginal utility of consumption. This equation is known as the *Consumption Capital Asset Pricing Model* or Consumption CAPM, and it plays an important role in modern finance. Most asset returns depend on payments generated by the real economy and so they are procyclical—they are better in expansions than during recessions. However, the usual assumption of diminishing marginal utility implies that U' depends negatively on consumption. This means that the covariance term is negative for assets whose returns are positively correlated with consumption and these assets will have a higher rate of return than the risk free rate. Indeed, the higher the correlation of the asset return with consumption, the higher will be the expected return.

Underlying this behaviour is the fact that consumers would like to use assets to hedge against consumption variations. Given two assets that have the same rate of return, a risk-averse consumer would prefer to have one that was negatively correlated with consumption than one that is positively correlated with consumption. For investors to be induced into holding both assets, the rate of return on the asset with a positive correlation with consumption needs to be higher.

Puzzles: Equity Premium and Risk-Free Rate

In theory, the consumption CAPM should be able to explain to us why some assets, such as stocks, tend to have such high returns while others, such as government bonds, have such low returns. However, it turns out that it has great difficulty in doing so. In the US, the average real return on stocks over the long run has been about six percent per year while the average return on Treasury bonds has been about one percent per year. In theory, this could be explained by the positive correlation between stock returns and consumption. In practice, this is not so easy. Most studies use simple utility functions such as the Constant Relative Risk Aversion (CRRA) preferences

$$U(C_t) = \frac{1}{1-\theta} C_t^{1-\theta} \tag{59}$$

so marginal utility is

$$U'(C_t) = C_t^{-\theta} \tag{60}$$

In this case, the consumption-CAPM equation becomes

$$E_t\left(R_{i,t+1}\right) = R_{f,t+1} - \frac{Cov\left(R_{i,t+1}, C_{t+1}^{-\theta}\right)}{E_t\left[C_{t+1}^{-\theta}\right]} \tag{61}$$

For values of θ considered consistent with standard estimates of risk aversion, this covariance on the right-hand side is not nearly big enough to justify the observed equity premium. It requires values such as $\theta = 25$, which turns out to imply people are incredibly risk averse: For instance, it implies they are indifferent between a certain 17 percent decline in consumption and 50-50 risk of either no decline or a 20 percent decline. One way to explain this finding is as follows. In practice, consumption tends to be quite smooth over the business cycle (our earlier model helps to explain why) so for standard values of θ, marginal utility doesn't change that much over the cycle and one doesn't need to worry too much equities being procyclical.

However, if θ is very very high, then the gap between marginal utility in booms and recessions is much bigger: Marginal utility is really high in recessions and consumers really want an asset that pays off then. This leads to a high equity premium.

One route that doesn't seem to work is arguing that people really are that risk averse, i.e. that $\theta = 25$ somehow is a good value. The reason for this is that this value of θ would imply a much higher risk-free rate than we actually see. Plugging the CRRA utility function into the equation for the risk free rate

$$R_{f,t+1} = \frac{(1+\beta)\, C_t^{-\theta}}{E_t\left[C_{t+1}^{-\theta}\right]} \tag{62}$$

Neglecting uncertainty about consumption growth, this formula implies that on average, the risk-free rate should be

$$R_f = (1+\beta)\,(1+g_C)^{\theta} \tag{63}$$

where g_C is the growth rate of consumption. Plugging in the average growth rate of consumption, a value of $\theta = 25$ would imply a far higher risk-free rate than we actually see on government bonds.

There is now a very large literature dedicated to solving the equity premium and risk-free rate puzzles, but as of yet there is no agreed best solution.[8]

[8] The paper that started this whole literature is Rajnish Mehra and Edward Prescott, "The Equity Premium: A Puzzle" *Journal of Monetary Economics*, 15, 145-161. For a review, see Narayana Kocherlakota, "The Equity Premium: It's Still a Puzzle" *Journal of Economic Literature*, 34, 42-71.

Things to Understand from these Notes

Here's a brief summary of the things that you need to understand from these notes.

1. The household budget constraint.
2. How to derive the intertemporal budget constraint.
3. How to set up and derive first-order conditions for optimal consumption.
4. How to derive the Rational Expectations/Permanent Income Hypothesis.
5. The Lucas Critique applied to temporary tax cuts.
6. The Ricardian equivalence hypothesis.
7. Evidence on temporary tax cuts.
8. Precautionary savings.
9. The first-order condition with time-varying asset returns.
10. The Consumption-CAPM model.
11. The equity premium and risk-free rate puzzles