already; they cant get collateral; they have criminal records, and so on.[5] And those who are *liquidity constrained* are stuck with the income they have. For this reason, with liquidity constraints, we can have $u'(C_1) > u'(C_2)$ and consumption tracking income. Consider the usual two-period world. With a constraint of $S_1 \geq 0$ and when $Y_1 < \frac{Y_1+Y_2}{2}$, we must have $C_1 = Y_1$ and $C_2 = Y_2$. However, if $Y_1 > \frac{Y_1+Y_2}{2}$, the consumer does not wish to borrow anyway, so the liquidity constraint doesn't matter (formally, we say the constraint doesn't *bind* in this case.) With liquidity constraints, the consumption function in the first period is $C_1 = \min\{Y_1, \frac{Y_1+Y_2}{2}\}$.

1.1 Multiperiod Version

Of course, in reality people live for many periods. In fact, it is common in macroeconomics to assume people are *infinitely lived*; namely, people live through their children and transfer wealth intergenerationally via bequests. Happily for us, since this rule holds for any two arbitrary periods, it holds for arbitrarily many periods too. Consider first what happens in the case of three periods. The utility function is $u(C_1) + u(C_2) + u(C_3)$ and the budget constraints are:

$$Y_1 = C_1 + S_1$$

$$Y_2 + S_1 = C_2 + S_2$$

$$Y_3 + S_2 = C_3$$

In accord with the transversality condition, there are no savings in the last period; i.e., $S_3 = 0$. Combining these conditions—just eliminate all the S terms—gives $Y_1 + Y_2 + Y_3 = C_1+C_2+C_3$. Then the first order conditions are $u'(C_1) = u'(C_2) = u'(C_3) \Rightarrow C_1 = C_2 = C_3$. And then the solution is $C_1 = C_2 = C_3 = \frac{Y_1+Y_2+Y_3}{3}$.

More generally, if you live for $T > 3$ periods, then the consumer's problem:

[5] Liquidity constraints are often a result of adverse selection and moral hazard issues. In the case of adverse selection, banks don't raise interest rates too high, since high rates attract risky borrowers—or "lemons"—who are unlikely to repay. Namely, borrowers who take out loans at high rights might do so, thinking they mightn't pay it back; for this reason, high rates might attract disproportionately risky borrowers. Instead of raising rates, they just deny credit to some borrowers. Meanwhile, with moral hazard, banks may be reluctant to lend anyone too much—"credit limits"—in case borrowers spend the money recklessly, in which case they might default.

$$\max_{\{C_t\}_{t=0}^{t=T}} \sum_{t=0}^{t=T} u(C_t) \quad \text{subject to} \quad \sum_{t=0}^{t=T} C_t = \sum_{t=0}^{t=T} Y_t$$

The Lagrangian is

$$\sum_{t=0}^{t=T} u(C_t) + \lambda \left(\sum_{t=0}^{t=T} Y_t - \sum_{t=0}^{t=T} C_t \right)$$

The solution now is:

$$C_1 = \frac{\sum_{i=1}^{T} Y_i}{T} = \ldots = C_T$$

That is, consumption again equals permanent income. And, finally if you will receive assets, A, at some point, then:

$$C_1 = \frac{A + \sum_{i=1}^{T} Y_i}{T} = \ldots = C_T$$

1.2 Stabilisation Policy

Having presented the basic idea, I now turn to some applications. What are the implications for fiscal policy? To see this, imagine you get a tax break of τ this period, thereby raising current income to $Y_1 + \tau$. Following the analysis above, our consumption *each period* is reduced to:

$$C_1 = C_2 = \frac{Y_1 + \tau + Y_2}{2},$$

and hence consumption today increases by only $\frac{\tau}{2}$. Therefore, *according to the PIH temporary government policies will have little power to stimulate the economy.* To see this formally:

$$\frac{\partial C_1}{\partial \tau} = \frac{1}{2}$$

And if—as in reality—you live for T periods:

$$\frac{\partial C_1}{\partial \tau} = \frac{1}{T}$$

Taking limits gives:

$$\lim_{T \to \infty} \frac{\partial C_1}{\partial \tau} = 0,$$

that is, as consumers' lifetimes increases, the stimulus becomes less and less effective.

Continuing this example, with a permanent tax cut of τ—giving an income stream of $Y_1 + \tau$ and $Y_2 + \tau$—we have:

$$C_1 = \frac{Y_1 + \tau + Y_2 + \tau}{2}$$
$$\Rightarrow \frac{\partial C_1}{\partial \tau} = 1$$

Thus, C_1 rises one for one, and you rationally spend all of a permanent change. This way, permanent changes in fiscal policy can have significant effects. But, almost by definition, stabilization policy is temporary! Except for the cases when tax breaks are permanent—mostly they're not—and people are liquidity constrained, stabilization policy is ineffective in theory. With binding liquidity constraints, people are hungry for money since they're not at their optima in the first place—so they'll dutifully spend what they get.[6]

However, largely because of the PIH, economists are skeptical of the power of fiscal policy, and as a result, regard monetary policy as the prime tool to stimulate an economy. Even more striking is what happens when rational consumers take account of the government's intertemporal budget constraint (more on this later).

Recall that the basic Keynesian multiplier was $\frac{1}{1-mpc}$.[7] The role of the multiplier was central to the IS-LM and Keynesian cross analysis . But for temporary income changes—like those in stabilisation policy—the PIH predicts the multiplier is very small. So if you think about it, the PIH has large implications for Keynesian economics: a small multiplier effect undermines much of its original appeal. Indeed, all of the current debates on fiscal policy are essentially debates on whether the PIH is correct.

1.2.1 A Note on Interest Rates

A quick word about interest rates. The *real* ex post rate of return on something is given by the equation:

$$r = i - \pi.$$

This indicates the real, *purchasing power* return on my investment.[8] *And this is all I care about.* This equation just captures the idea that inflation "eats away" at nominal returns. Just think of the real interest rate as a measure of how many *goods* you get back (as I said,

[6]Having said this, if they expect to be liquidity constrained in the future too, they'll save some to be *less* liquidity constrained henceforth.

[7]To see this, recall that $Y = C + I + G$ which means $Y = c_0 + mpc\ Y + I + G$. This implies $Y = \frac{c_0+I+G}{1-mpc} \Rightarrow \frac{\partial Y}{\partial G} = \frac{1}{1-mpc}$.

[8]Strictly speaking, this is an approximation that is only valid for small levels of inflation.

it's much easier to think in terms of goods). For example, if $r = .05$, and if I lend you one good, I get 1.05 goods in return.

Now, lots of investors lost out in the $70's$ since they bought bonds and only after did inflation rear its ugly head. This diminished their *real* returns. In other words, they lent out money but the purchasing power of what they got back was much less. For instance, if I lent \$10 to someone a hundred years ago and I got a mere \$12 back today, then, despite a 20% nominal return, this has hardly any *purchasing value compared* to what I lent out, given the enormous price level increase in the interim. Ex post, then, inflation is the borrower's friend, since it reduces the real rate of interest or real burden of payment.

Speaking of which, what do I mean when I say you offered me a rate of interest, i? Doesn't the central bank—say, Bernanke—control i? Well, not really. Bernanke controls what we call the *federal funds rate*: the rate at which banks lend to each other (so as to satisfy their reserve requirements stipulated by the FED). But more important is the role of *long-run interest rates*, which are set by market forces in financial markets. However, the federal funds rate and all other interest rates generally *move together*. If the banks have to pay more on loans from other banks, they'll dutifully pass that on to customers in the *prime rate*. And if the interest rates in the banks are high, then corporate bonds will have to pay a higher return too. Bottom line is that all rates tend to move together. Because all interest rates move together and we are only concerned with changes in interest rate, for now I will refer to just "the interest rate."

1.3 Interest rates and Intertemporal Choice

Which brings us to the next topic. Up until now, we have assumed away issues with interest and discount rates. Although the main insights remain intact, it is interesting to ask: *Under what circumstances, do we deviate from perfect smoothing (assuming certainty)?* Well, there are two ways: *Either we prefer the present or we are rewarded from postponing consumption.* Interest rates are a way to lure or seduce investors from perfect consumption smoothing; this will tend to *increase* future consumption. Meanwhile, a low discount factor (β)—i.e., a high rate of time preference—means you get more utility from consuming today; in contrast, this will tend to *decrease* future consumption. But just to be clear: these issues are do not overturn the main idea of consumption smoothing. One more thing: In this partial equilibrium part of the course, we assume consumers take the interest rate as given.[9]

[9] In a general equilibrium setting, the interest rate is endogenous: it would change along with the level of savings. In addition, to compensate for risk of default, the interest rate is often a function of the level

First, I'll derive the optimal conditions with Lagrangians and then present two other ways.

Case when $r \neq 0$ and $\beta \neq 1$.

With these additional frills, utility is now:

$$\max_{C_1 \geq 0, C_2 \geq 0} u(C_1) + \beta u(C_2); \qquad \beta \in [0,1]. \tag{1.3}$$

The budget constraints for period one and two are:

$$C_1 + S = Y_1$$

$$C_2 = (1+r)S + Y_2$$

Plugging the first into the second:

$$C_2 = (1+r)(Y_1 - C_1) + Y_2$$

And manipulating this gives:

$$\underbrace{C_1 + \frac{C_2}{1+r}}_{uses} = \underbrace{Y_1 + \frac{Y_2}{1+r}}_{sources}$$

After doing all this, the consumer's problem reduces to:

$$\max_{C_1 \geq 0, C_2 \geq 0} U(C_1, C_2) = u(C_1) + \beta u(C_1),$$

subject to:

$$C_1 + \frac{C_2}{1+r} = Y_1 + \frac{Y_2}{1+r}$$

Setting up the Lagrangian gives:

$$\mathbb{L} = u(C_1) + \beta u(C_2) + \lambda (Y_1 + \frac{Y_2}{1+r} - C_1 - \frac{C_2}{1+r})$$

Then taking first order conditions with respect to C_1 and C_2 gives:

$$u'(C_1) = \lambda$$

of borrowing itself. For instance, because of increased borrowing, the Irish government must now pay a substantially higher interest rate when it borrows.

$$\beta u'(C_2) = \frac{\lambda}{1+r}$$

Combining:

$$\boxed{u'(C_1) = \beta(1+r)u'(C_2)}$$

This is the *Euler Equation*. Implicitly, this condition pins down the optimal path of consumption. As before, to find the optimal *level* of C_1 and C_2, we must combine this with the intertemporal budget constraint.

For instance if $r = 0$, we have:

$$u'(C_1) = \beta u'(C_2) \Rightarrow u'(C_1) < u'(C_2) \Rightarrow C_1 > C_2.$$

The reason $C_1 > C_2$? Consumers derive more utility from consumption in period 1; hence the bias their consumption profile towards the first period. The opposite effect happens for a positive interest rate, $r > 0$: consumption will rise over time.[10] So, except for the case where $(1+r)\beta = 1$, we no longer have perfect consumption smoothing. If $\beta(1+r) = 1$, then we are—quite naturally—back to the same situation as before. In summary, the trajectory of consumption over time depends on the "tug of war" between r and β.

1.3.1 Alternative Ways of Deriving Euler Equation

Conversion into One-Variable Problem

$$\max_{C_1 \geq 0, C_2 \geq 0} u(C_1) + \beta u(C_2)$$

Substituting

$$C_2 = (1+r)(Y_1 - C_1) + Y_2$$

into $u(C_1) + \beta u(C_2)$ gives

$$u(C_1) + \beta u((1+r)(Y_1 - C_1) + Y_2).$$

Then, maximizing the above with respect to C_1 (and noting the chain rule) gives

[10] Yet this only tells us that there will be positive consumption growth. It does not tell us whether consumption falls in period 1 or not. For example, we could start off with $\beta = 1$, $r = 0$ and $C_1 = C_2 = 10$. With a positive r, we would then have $C_1 < C_2$. But this could hold true even if $C_1 = 11$ and $C_2 = 13$ or when $C_1 = 9$ and $C_2 = 14$.

$$u'(C_1) - (1+r)\beta u((1+r)(Y_1 - C_1) + Y_2) = 0$$

Then substituting back in $C_2 = (1+r)(Y_1 - C_1) + Y_2$ gives

$$u'(C_1) - \beta(1+r)u(C_2) = 0 \Rightarrow u'(C_1) = \beta(1+r)u(C_2).$$

Law of Equi-marginal Utility

The interest rate is the relative price of future consumption. Why? A high interest rate makes future consumption cheaper. Because you give up one unit today and receive more tomorrow in exchange, future units—i.e., future consumption—are now effectively cheaper. Formally, the relative price of consumption in period 2 is $\frac{1}{1+r}$. Now remember the law of equimarginal returns—i.e., $\frac{MU_i}{p_i} = \frac{MU_j}{p_j}$ for all goods i and j— where you equated the "bang per buck" across goods? One can view the Euler equation as a special case thereof, where the "goods" refer to consumption in each period. Using this condition, equilibrium quantities are then implicitly defined by:

$$u'(C_1) = \frac{\beta u'(C_2)}{\frac{1}{1+r}}$$

Of course this is just our friend again.

Arbitrage

Suppose we are at the optimum C_1 and C_2. Then the *marginal loss* from reducing C_1 by one unit is $u'(C_1)$. Note that we get back $1+r$ units which provide a utility of $u'(C_2)$. And since next periods utility is discounted by β, the *marginal benefit* is $\beta(1+r)u'(C_2)$. So, overall:

$$\begin{aligned} u'(C_1) &\quad \ldots \text{marginal cost} \\ \beta(1+r)u'(C_2) &\quad \ldots \text{marginal benefit.} \end{aligned}$$

Now *since we were at an optimum*, the net gain to this change must be zero (else, it wouldn't have been an optimum!) Hence:

$$-u'(C_1) + \beta(1+r)u'(C_2) = 0 \Rightarrow u'(C_1) = \beta(1+r)u'(C_2)$$

1.3.2 Multiperiod Version

With many periods, consumers solve:

$$\max_{\{C_t\}_{t=0}^{t=T}} \sum_{t=0}^{t=T} \beta^t u(C_t) \quad \text{subject to} \quad \sum_{t=0}^{t=T} \frac{C_t}{(1+r)^t} = \sum_{t=0}^{t=T} \frac{Y_t}{(1+r)^t}$$

To solve this, we again use the Lagrangian technique. Assuming interest rates are constant over time, the Lagrangian is

$$L = \sum_{t=0}^{t=T} \beta^t u(C_t) + \lambda \sum_{t=0}^{t=T} \left(\frac{Y_t}{(1+r)^t} - \frac{C_t}{(1+r)^t} \right).$$

And with initial assets of A, this would be

$$L = \sum_{t=0}^{t=T} \beta^t u(C_t) + \lambda \sum_{t=0}^{t=T} \left(A + \frac{Y_t}{(1+r)^t} - \frac{C_t}{(1+r)^t} \right).$$

Implicit in the intertemporal constraint is the TVC.

Depending on whether the interest rate or discount rate force dominates, consumption will either rise or fall over time. Solving this would yield a set of Euler equations: $u'(C_1) = \beta(1+r)u'(C_2)$, $u'(C_2) = \beta(1+r)u'(C_3)$, $u'(C_3) = \beta(1+r)u'(C_4)$, etc. Note how this implies $u'(C_1) = \beta^3(1+r)^3 u'(C_4)$, and if interest rates were different, we'd have $u'(C_1) = \beta^3(1+r_1)(1+r_2)(1+r_3)u'(C_4)$; that is, consumption today depends on the path of future interest rates. This way, we can relate consumption today to consumption far off in the future and long-run interest rates.

1.3.3 Functional Form for Utility

So far, we have just derived an expression for the growth of marginal utility. Still, we haven't found the optimum *levels* of C_1 and C_2. Unlike the first case, we cannot simply average income over time. But, considering both the Euler equation and budget constraint, we now have two equations in two unknowns, C_1 and C_2. To solve for levels, we must posit a functional form for utility.

The most common utility function in macroeconomics takes the form:

$$u(C) = \frac{C^{1-\theta}}{1-\theta}, \qquad \theta > 0$$

This implies marginal utility is

$$u'(C) = C^{-\theta} = \frac{1}{C^\theta}.$$

Note that the higher θ is, the more quickly DMU sets in. Moreover, it's strictly concave since:

$$u''(C) = -\frac{\theta}{C^{\theta+1}} < 0.$$

With this function, lifetime utility is:

$$U(C_1, C_2) = \frac{C_1^{1-\theta}}{1-\theta} + \beta\frac{C_2^{1-\theta}}{1-\theta}.$$

Let's talk about this for a moment. Consider θ. This parameter tells us how quickly DMU sets in; to be specific, θ is the percentage fall in marginal utility when consumption rises by one percent. Overall, it measures the curvature of the utility function. Graphically, a utility function with a high θ flattens out quickly.[11]

Remember, you are concerned about the utility gain from shifting consumption around. That's all that matters. If DMU sets in really quickly, it makes no sense to have lot of consumption in any *given* period. With DMU, what's the point? Consider this: Instead of having a lunch today and tomorrow, would you rather have two lunches today? Well, no. Given DMU to lunch sets in pretty quickly, you *aggressively* try to smooth out lunch consumption. And this level of aggressiveness has a name: *the intertemporal elasticity of substitution*, which is mathematically given by $\frac{1}{\theta}$. Thus, if DMU sets in really quickly—i.e., θ is high—your intertemporal elasticity of substitution is low. Because responding to interest rates involves shifting consumption forward, this parameter measures how responsive consumers are to changes in interest rates.

To see what I'm talking about, consider two *goods*: salt and luxury yachts. For a good like salt, people want to consume only a little each day. In particular, they don't want too much salt in one period and none in the other (you see, food is tasteless without salt.) In other words, there is sharply diminishing marginal utility to salt. As a result, the IES for salt is likely very low. If all goods were like salt, would people increase reducing consumption and savings in response to a higher interest rate. I doubt it. That means we'd have little salt this period and lots next period—hardly an attractive option. By contrast, consider the luxury yachts. Realistically, you could do without a yacht this period and have one tomorrow instead. So for a good like this—that's not essential—consumers would be more willing to shift them around; formally, the IES for this good would be relatively high. The overall IES for consumption depends of course on whether the average good is more like salt

[11]Notice that if $\theta > 1$ this function is negative. Since utility is only used to compare things, this is just fine. In this setting, if utility becomes less negative, there's a welfare improvement; that's all we're interested in.

or yachts. The fact that the IES is low empirically suggests the average good is rather like salt.[12]

As noted, θ governs how willing you are to shift consumption around. When we incorporate uncertainty into models, the parameter θ is called the coefficient of relative risk aversion. It measures risk since risk entails the basic idea of valuing losses and gains. You see, if DMU sets in really quickly (i.e., θ is high), then gains are basically worthless in terms of marginal utility. Meanwhile, losses are still painful. Empirically, θ is often measured by looking at people's choices in risky situations. For instance, what the wage premia for risky occupations?

Interest Rates

Assume now you are deciding how much of this period's income to save. From now on, I am also assuming the consumer is a *saver* in period 1. How does a rise in the interest rate—say a doubling—affect your plans, in particular today's consumption? Whether C_1 rises or falls (relative to the previous optimal plan) upon a rise in interest rates depends in part on the interaction of income and substitution effects. But, as we shall see, there are three effects. First, there is the *substitution effect*; as with all substitution effects, it deals with the change in relative prices. Now that returns to saving are higher, you should "make hay while the sun shines" and therefore save more. Put another way, a rise in the interest rate makes today's consumption relatively more costly. And this makes you consume *less* today. In short, the substitution effect says: *go for it, save more.*

Second, there is the *income effect*: now you can attain a *given* level of savings (i.e., future consumption) with less work, so you are effectively richer. Equivalently, you are richer, since the price of future consumption is now cheaper. And seeing you are now richer, there's less need for saving; you should consume *more* today (*and* next period). So the income effect says: *look, you're now better off; save less.*

So now what? Depending on the strengths of the income and substitution effects, consumption can clearly go either way in period 1. Which effect is stronger?[13] Happily for us, though, we *can* actually tell which effect is stronger from the consumer's utility function; in particular, from the *intertemporal elasticity of substitution.* Because this tells us how extra

[12] One could rationalize this by saying consumers become attached to different goods over time. For instance, 10 years ago, most people could have done without the internet. Yet, today, the internet has become virtually essential—making it like salt, so to speak.

[13] For the second period, however, income and substitution effects go in the same direction. Note that the substitution effect dictates more consumption in period 2 due to the lower relative price. The income effect dictates more consumption in *both* periods.

units of consumption are valued in a *given* period, it naturally governs the consumer's desire to shift consumption across periods. In turn, in this example, it governs the consumer's response to interest rate changes, and specifically the substitution effect (i.e., how willing is the consumer to "transfer" consumption from this period to the next?). Turns out, if $\theta < 1$, the substitution effect will dominate the income effect. And it's the other way round for $\theta > 1$; and of course effects just balance if $\theta = 1$.

However, there is a *third* effect. A rise in the interest rate reduces the present discounted value of lifetime income, $Y_1 + \frac{Y_2}{1+r}$. To see why, recall that the present discounted value gives the value today of what I get in the future, Y_1. Equivalently, it answers the question: what do I have to invest today to get my future income, Y_1. Therefore, with a large interest rate, my future income is worth less today; namely, if the interest rate is larger I only need a small amount today to get a given amount, Y_1, in the future; as a result, my future claim is worth less in today's terms. In this sense, a higher interest rate reduces the today's value of future income and makes the consumer feel poorer. Because of this, a higher interest rate works attenuates the income effect, making it more likely that the substitution effect will dominate. Of course, the magnitude of this depends on how much income one has in the future; if $Y_2 = 0$, this effect is absent. For a younger person, therefore, this effect would be larger.

An Example

Euler Equation with CRRA Utility

$$U(C_1, C_2) = \frac{C_1^{1-\theta}}{1-\theta} + \beta \frac{C_2^{1-\theta}}{1-\theta}$$

subject to

$$Y_1 + \frac{Y_2}{1+r} = C_1 + \frac{C_2}{1+r}$$

The Euler equation, $u'(C_t) = \beta(1+r)u'(C_{t+1})$, reduces to (setting $\beta = \frac{1}{1+\rho}$):

$$\frac{C_2}{C_1} = (\beta(1+r))^{\frac{1}{\theta}} \Rightarrow C_2 = C_1 (\beta(1+r))^{\frac{1}{\theta}}$$

Combining with the intertemporal budget constraint gives

$$C_1 = \left(Y_1 + \frac{Y_2}{1+r}\right) \frac{1}{1 + (1+r)^{\frac{1}{\theta}-1} \beta^{\frac{1}{\theta}}} \tag{1.4}$$

The most important point to note here is that consumption depends on the present discounted value of lifetime. As well, if the consumer attained assets of A in period two, say, then C_1 would become: $C_1 = \left(Y_1 + \frac{Y_2+A_2}{1+r}\right)\frac{1}{1+(1+r)^{\frac{1}{\theta}-1}\beta^{\frac{1}{\theta}}}$. Note, too, that in the case of $r = 0$ and $\beta = 1$, we get the familiar result, $C_1 = \frac{Y_1+Y_2}{2}$. The marginal propensity to consume out out Y_1 is

$$\frac{\partial C_1}{\partial Y_1} = \frac{1}{1+(1+r)^{\frac{1}{\theta}-1}\beta^{\frac{1}{\theta}}}.$$

Observe that $\frac{\partial C_1}{\partial \beta} < 0$, while the sign of $\frac{\partial C_1}{\partial r}$ is indeterminate (without knowing θ.) And because $S_1 = Y_1 - C_1$, we have $\frac{\partial S_1}{\partial r} = -\frac{\partial C_1}{\partial r}$, so the change in savings is opposite to the change in consumption (since they are two sides of the same coin.)

Permanent and Temporary Changes

So far I have implicitly assumed changes in interest rates were permanent. By construction, this had to be the case in a two-period world. However, whether a change in interest rates is temporary or permanent matters a lot in a multi-period world. To see why, suppose a consumer lives for fifty periods, and the interest rate rises *temporarily* in period one. (Interest rates will revert to normal again in period 2.) Assume further the consumer receives all income in period 1. What happens? Well, consider the income and substitution effects. The substitution effect dictates the consumer should save more. The income effect says the consumer is richer and should save less. But—and here's the but—the income effect is relatively weak in this situation. Namely, since the consumer lives for fifty periods and the interest rate rises only for one period, the consumer doesn't feel *that* much richer as a result. Sure, he gets more interest on *this* period's savings, but, alas, he lives for fifty periods. But it's the presented discounted value of all income that matters to him, and this is changed relatively little. It should be clear that the income effect is smaller than in the case where the interest rate rises *permanently.* By contrast, the strength of the substitution effect remains the same. As a result, in the case of a *temporary* interest rate rise, the substitution effect will likely dominate, and the consumer will respond by raising savings.

1.4 Uncertainty

In the case where the return and consumption are uncertain—most returns *are* uncertain—we must put an expectations operator on the right hand side. A prominent source of consumption uncertainty arises from a stochastic income stream. In this case, however, we typically know the *distribution* of consumption i.e., its statistical moments. So $u'(C_1) = \beta(1+r_i)u'(C_2)$ becomes $u'(C_1) = \mathbb{E}_1\beta(1+r_i)u'(C_2)$, when the payoff r_i is uncertain (formally: when the payoff, r_i, is a random variable.) The notation $\mathbb{E}_t$ indicates *expectation at period t*.

1.4.1 The Stochastic Euler Equation

From the standpoint of time t

$$u(c_t) + \mathbb{E}_t\beta u(c_{t+1})$$

Writing in terms of savings s

$$u(y_t - s_t) + \mathbb{E}_t\beta u((1+r)s_t)$$

Differentiating w.r.t s_t gives

$$-u'(c_t) + \mathbb{E}_t\beta(1+r)u'(c_{t+1}) = 0$$

$$u'(c_t) = \mathbb{E}_t\beta(1+r)u'(c_{t+1})$$

In this case, we can never nail down a consumption path *ex ante*. What we get are a series of rules—*contingent plans*—that relate periods to each other. For instance, suppose we have three periods and $\beta(1+r) = 1$. The uncertainty relates to income uncertainty. Then, from the standpoint of period 0

$$u'(c_0) = \mathbb{E}_0 u'(c_1) = \mathbb{E}_0 u'(c_2)$$

That is, we set expected marginal utilities equal each period. Because of the uncertainty, we can never ascertain the precise consumption path *ex ante*. However, the plan will change as more information arrives. In this example, in period 1, when the period 1 uncertainty has been revealed, you will set

$$u'(c_1) = \mathbb{E}_1 u'(c_2)$$

More concretely, we say there is a series of stochastic Euler equations for this problem; namely

$$u'(c_0) = \mathbb{E}_0 u'(c_1)$$

and

$$u'(c_1) = \mathbb{E}_1 u'(c_2)$$

Let's consider a numerical example. Suppose income in period 0 is 5, while *expected* income is 6 in period 2 and 10 in period 3. Looking forward from the standpoint of period 0, you plan to consume 7 each period. So you consume 7 in period 1. But then period 2 arrives, and, just for the fun of it, suppose income is 100. Are you going stick to your initial plan? Of course not. Now expected lifetime income for the remaining periods is $100 + 10 = 55$. So you'll plan to consume 55 in both. Formally, you set $u'(c_1) = \mathbb{E}_1 u'(c_2)$. This way, the consumption plan changes as you receive more information. [14]

1.5 Precautionary Savings

From the stochastic Euler equation with $(\beta(1+r) = 1)$, we have (for periods 1 and 2)

$$u'(c_1) = \mathbb{E}_t u'(c_2)$$

$$u'(y_1 - s) = \mathbb{E}_t u'((1+r)(y_1 - s) + \epsilon)$$

where ϵ is the source of the uncertainty and $\mathbb{E}_t \epsilon = 0$, but $\sigma_\epsilon^2 \neq 0$. So instead of receiving $(1+r)(y_1 - s)$ next period, the consumer will receive $(1+r)(y_1 - s) + \epsilon$. Because $\mathbb{E}_t \epsilon = 0$, the consumer still expects to receive $(1+r)(y_1 - s)$. For instance, the consumer could receive either $(1+r)(y_1 - s) + 100$ with probability $\frac{1}{2}$ or $(1+r)(y_1 - s) - 100$ with probability $\frac{1}{2}$; this way, expected income is still $(1+r)(y_1 - s)$. Precautionary savings arise when the consumer is particularly fearful of receiving the bad draw, $(1+r)(y_1 - s) - 100$.

To start with, consider a simple numerical example. Suppose ϵ can take value of either 1 or -1; this way, $\mathbb{E}_t \epsilon = 0$. Now consider the *convex function*, $f(x) = x^2$.[15] Therefore, in period 1, we have $\mathbb{E}_t f(x) = \frac{1}{2}(1)^2 + \frac{1}{2}(-1)^2 = 1$. Suppose now ϵ can take value of either 10 or -10. In this case, $\mathbb{E}_t f(x) = \frac{1}{2}(10)^2 + \frac{1}{2}(-10)^2 = 100$. The central point here is that greater

[14] As we will see below, this will not necessarily be the case the consumption plan will evolve this way. This particular example demonstrates the case of *certainty equivalence*; i.e., where consumers do not engage in precautionary savings.

[15] Technically, a convex function is one where a line joining any two points on the curve lies *above* the curve. Intuitively, a convex function is one which gets disproportionately large in magnitude as we move in certain directions.

variance in ϵ raises $\mathbb{E}_t f(x)$ when f is a strictly *convex* function.[16] Following on from this, if the u' is a convex function and the ϵ in $\mathbb{E}_t u'((1+r)(y_1 - s) + \epsilon)$ becomes more variable, $\mathbb{E}_t u'(c_{t+1})$ will rise.[17] According to the Euler equation, $u'(c_t)$ also rises; i.e., consumption in period t falls and savings rises. For this reason, we say the rise in uncertainty induces a rise in *precautionary* savings.

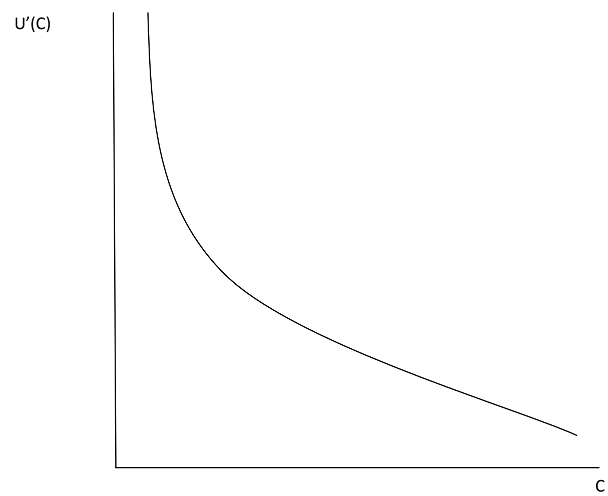

Figure 1.1: If marginal utility takes this form, then precautionary savings will arise. Intuitively, low levels of consumption lead to disproportionately high levels of marginal utility. Because of this, consumers will fear "bad draws" and will save more.

[16]Note that this is not the case if f is linear, for instance. Say $f(x) = x$. Now, when ϵ can take values of 1 or -1, $\mathbb{E}_t f(x) = \frac{1}{2}(1) + \frac{1}{2}(-1) = 0$. And when ϵ can take on values of -10 or 10 with equal likelihood, then $\mathbb{E}_t f(x) = \frac{1}{2}(10) + \frac{1}{2}(-10) = 0$; i.e., variance does not matter for $\mathbb{E}_t f(x)$.

[17]This property of marginal utility (shown graphically in Figure 1.1) is a feature of all standard utility functions used in macroeconomics.