EC 3010: Macroeconomics

October 6, 2010

Chapter 1

The Permanent Income Hypothesis

Developed in the 1950s by Milton Friedman, the permanent income hypothesis is now the standard theory of consumption in macroeconomics. Yet, to fully appreciate its implications, we must go back to the standard Keynesian consumption function:

$$C = c_0 + mpc\, Y, \quad c_0 > 0; \quad 0 < mpc < 1,$$

that is, consumption is a fixed fraction, *mpc*, of income plus come constant, c_0—autonomous consumer spending. Although this function is appealing for its simplicity, it has three main shortcomings. First, and most importantly, the function fails to distinguish between temporary and permanent income. To take an example, say you win a prize of 1000 euros, and so your income this period increases. Now ask yourself: would your consumption response be the same if you received a *permanent* income rise of 1000 (due say to promotion)? According to the Keynesian consumption function, it is: it treats all income equally. But this does not seem right: most likely, people would respond more to a permanent change in income; after all, the change is, well, permanent. Second, the function is not based on microeconomic foundations. That is, this function is not derived from a consumer's maximization problem. Rather, it was simply *assumed* by Keynes in what he called a "psychological law."

Third, and finally, its implication for the consumption-income ratio is counterfactual. To see why, divide across by Y to get:

$$\frac{C}{Y} = \frac{c_0}{Y} + mpc.$$

Because $\frac{c_0}{Y}$ falls as income, Y, increases, the Keynesian consumption function predicts—counterfactually—that the consumption income ratio, $\frac{C}{Y}$, falls as an economy grows. Yet in reality, $\frac{C}{Y}$ is approximately constant as GDP increases over time; hence, C and Y grow at the same rate.

Now, to incorporate the future into our analysis and differentiate between temporary and permanent in a rigorous way, we turn to the *permanent income hypothesis* (PIH). Essentially this states that consumption should depend on normal or *permanent* income, where *permanent income* is a function of the present discounted value of *all* lifetime income —basically, it is the average income a person expects to have over their lifetime; for example, if my only income is 100 euros in ten years time, my *permanent income* each year is 10 (ignoring interest rates). As a result, savings will be high when disposable income is higher than permanent income, and conversely. This seems obvious, and it is. Less obvious, however, are its implications. But before we turn to them, we must show where the PIH comes from.

Diminishing Marginal Utility

First, some background. Underlying the permanent income hypothesis is the basic idea of diminishing marginal utility to consumption (DMU). Just imagine listening to a song or eating food: the last unit of consumption is never quite as good as the first. Point is, the "bang per buck" falls as consumption of a good rises.[1] For this reason, we assume the utility function is strictly concave; that is, its derivative with respect to consumption—i.e., marginal utility or "the bang per buck"—falls as consumption rises. This way, the idea of concavity captures the realistic notion of DMU.

Assuming the utility function is $u(C)$, marginal utility is:

$$u'(C) > 0.$$

Think of marginal utility as a measure of how "hungry" you are for more consumption. Because we assume marginal utility is always positive ("more is better" or "nonsatiation"), consumers always want *more* and so ultimately consume *all* their income.

Moving on, our utility function $u : \mathbb{R} \to \mathbb{R}$ satisfies the *Inada conditions*:

- $$\lim_{C \to 0} u'(C) = \infty.$$

[1] Here, I'm assuming away anomalous cases such as addictive goods. With such goods, marginal utility might rise as consumption increases.

Interpret this: as consumption falls to zero you become extremely "hungry." Keep in mind that when consumption is low, marginal utility is high.

- $$\lim_{C \to \infty} u'(C) = 0.$$

  Interpret this: as consumption becomes arbitrarily large, the consumer becomes satiated. Unhappily for you, your marginal utility is a lot higher than that of Bill Gates.

Even though consumption, C, formally refers to real consumption expenditure in a period, from now on, it is useful to imagine you are living in a one-good world, where all goods have a price of one. So instead of thinking about the rather cumbersome "real consumption", just think about the *quantities* of some *given* good—say, coconuts—each period. This just makes things a little more intuitive.

Examples

For example, as a consequence of DMU, content yourself that:

$$u(4) + u(4) > u(3) + u(5)$$

And if utility is logarithmic (a fairly common case):

$$u(C) = \log C \Rightarrow u'(C) = \frac{1}{C} \Rightarrow u''(C) = -\frac{1}{C^2}.$$

Observe that this function satisfies the Inada conditions above. Nonetheless, convex—i.e., $u'' \geq 0$— and linear—i.e., $u'' = 0$—functions do not exhibit DMU.

Now, onto the derivation.

1.0.1 Derivation of The Permanent Income Hypothesis

To keep things simple, I will present the simplest case of two periods. Ignore uncertainty, and assume perfect capital markets; i.e., it's easy to get a loan. And for now, assume the interest rate, $r = 0$, and the discount factor, $\beta = \frac{1}{1+\rho} = 1$. Just to remind you, the parameter ρ is the rate of time preference; for example, a high ρ ($\Rightarrow$ low β) means you are impatient and place less weight on the future. But for the moment I assume you value the future as much as today; i.e., $\beta = 1$.

Furthermore, I assume the existence of a *representative agent*. By making this assumption, we are implicitly assuming little differences simply "wash out" in the aggregate. Consider, for instance, the EC 3010 class. Sure, some people have a high β; others have a low one. But there's someone who has the average class β: This is our representative consumer.

We want to maximize utility over two periods, say the *young* period and the *old* one. Or, "now" and "forever after." (Which reminds me, this analysis is completely analogous to how microeconomists analyze the spreading of consumption across *goods*). Now lifetime utility is $U(C_1, C_2)$, and the consumer solves:

$$\max_{C_1 \geq 0, C_2 \geq 0} U(C_1, C_2) = u(C_1) + u(C_2); \quad u'' < 0. \tag{1.1}$$

Of course, without the constraints, the solution is $C_1 = C_2 = \infty$. Unhappily, though, the consumer must obey the constraints in each period. Letting S denote savings and Y income, the constraints for period 1 and 2 are:

$$C_1 = Y_1 - S$$

$$C_2 = Y_2 + S$$

Clearly, then, consumption tomorrow is a function of postponed consumption today, S; thus, savings today is just consumption *tomorrow*. Note too that S can be negative, as in the case of borrowing. But instead of talking about savings, we can also write these constraints as

$$C_1 = Y_1 - B$$

$$C_2 = Y_2 + B,$$

where B denotes the quantity of bonds purchased (assuming the consumer saves via purchasing bonds.) Because period two is the last period, there are no savings in that period; instead, the consumer eats all the remaining wealth. This condition whereby the consumer does not leave any assets or debt leftover at the end is called a *transversality condition.*

Either way, combining these gives the *intertemporal budget constraint*:

$$C_1 + C_2 = Y_1 + Y_2,$$

where the income stream, Y_1 and Y_2, is given exogenously. Just to be clear, we are implicitly assuming consumers can borrow and lend easily (i.e., perfect capital markets);

that's why the budget constraint has lifetime income, $Y_1 + Y_2$, in it. But with borrowing constraints, we'd have to impose $C_1 \leq Y_1$; but ignore these for now. While we're talking about budget constraints, keep in mind that all budget constraints are of the form:

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

Now there are three ways to solve this constrained maximization problem. Because these methods are used throughout the course, this time I will present all three.

Unconstrained Maximization or "Direct Attack"

Probably the most familiar way is to simply turn the constrained maximization problem into an unconstrained one. By doing see, we change a two variable maximization problem into a one variable one. Start with the budget constraint, and isolate C_2 to get:

$$C_2 = Y_2 + Y_1 - C_1.$$

Plugging this into lifetime utility yields:

$$u(C_1) + u(Y_2 + Y_1 - C_1)$$

Then maximize with respect to C_1:

$$u'(C_1) - u'(C_2) = 0 \Rightarrow u'(C_1) = u'(C_2) \Rightarrow C_1 = C_2$$

Because of strict concavity, $u'' < 0$, these first order conditions are also sufficient for a maximum.

Method of Lagrangian Multipliers

Another common way is to use the Lagrangian technique. For this problem, the Lagrangian is:

$$\mathbb{L} = u(C_1) + u(C_2) + \lambda(Y_1 + Y_2 - C_1 - C_2). \tag{1.2}$$

Taking partial derivatives with respect to C_1, C_2, and λ gives:

$$\frac{\partial \mathbb{L}}{\partial C_1} = u'(C_1) - \lambda = 0 \Rightarrow u'(C_1) = \lambda$$
$$\frac{\partial \mathbb{L}}{\partial C_1} = u'(C_2) - \lambda = 0 \Rightarrow u'(C_2) = \lambda$$

$$\frac{\partial \mathbb{L}}{\partial \lambda} = 0 \Rightarrow Y_1 + Y_2 = C_1 - C_2$$

So combining we have:

$$u'(C_1) = u'(C_2) \Rightarrow C_1 = C_2.$$

Arbitrage

Think of marginal utility each period as the "return" to consuming in that period. To maximize total returns, then just equate returns across periods. Obviously, it makes no sense to invest in one "asset" (i.e., period), while ignoring another (period) with higher returns. Put a little differently, think of marginal utility as a measure of hunger. So why have different levels of hunger in both periods? Think about it: if we had a higher marginal utility (i.e., *lower* consumption) in one period than another, then that would be suboptimal, since we can increase lifetime utility by taking increasing consumption in that period and decreasing it in the other. And we should continue doing this until marginal utilities are equated, and no more welfare-enhancing "transferring" can occur.

More formally, suppose we are on an *optimal* consumption path; that is, C_1 and C_2 maximize lifetime utility. Suppose now I reduce consumption by a bit in period 1 and transfer it to period 2. The marginal cost and marginal benefit of this change are:

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

But given we were at an interior optimum, this change cannot increase utility; otherwise this would have been part of the optimal plan. But it's not. Considering that it was an optimal path (by definition), then we can't do any better. Therefore, the net utility change to this switching around must be zero:

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

So it's the *same old story*: Just equate marginal cost to marginal benefit. As a result, when C_1 and C_2 are the optimal plan, they must satisfy the condition, $C_1 = C_2$.

1.0.2 Discussion of Permanent Income Hypothesis

The conclusion? Clearly the problems all have the same solution: equate consumption over time.[2] Now for the best part: substituting the optimal condition, $C_1 = C_2$ into the intertemporal budget constraint gives the new *consumption function*:[3]

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

Even though the income stream might be volatile, the consumption profile is now *flat*. Moreover, C_1 and C_2 rise with lifetime income; they are normal goods. Compare this now to the Keynesian consumption function. In contrast, consumption now depends only on your *permanent income*, $\frac{Y_1+Y_2}{2}$. To take an example, consider this: say you expect to receive 100 over ten periods, implying 10 is your permanent income. For example, if you receive 15 in first period, you have *transitory income* of 5, which you will save. Rather than only considering today's income, you consider *all* future income when determining today's consumption; really, this is the very essence of the *permanent income hypothesis*. One important implication of this is that there should be no *expected* jumps in consumption.

Result 1 *According to the permanent income hypothesis, consumption only depends on present discounted value of lifetime income.*

What's going on here should be clear: precisely because of diminishing marginal utility to consumption in any *given* period, to maximize utility, consumers spread consumption as thinly as possible across *all* periods. Like the way people spread butter over bread, it's optimal to spread consumption over a number of periods. And for the same reason too: would it make sense to put too much butter on part of the sandwich, leaving other parts dry? Of course not; people optimally equate the "tastiness"—or marginal utility—of the sandwich at each part. Now, you might think people couldn't solve this maths problem. But, as with the sandwich, it's reasonable to assume people act *as if* they're solving a maximization problem. Of course, they don't literally do so; instead, they follow the Nike strategy and *Just Do It*.

[2] Sometimes, this result is called the lifecycle hypothesis. But for our purposes, there's really no difference between the PIH and lifecycle hypothesis. Technically, however, the PIH includes the case of living for infinitely many periods (i.e., through you offspring).

[3] Keep in mind that the first order conditions only tell us *relative consumption* in both periods; that is, how does C_1 compare to C_2? They do not tell us anything about actual consumption levels; these can be satisfied for $C_1 = C_2 = 10000$ and, say, $C_1 = C_2 = 0$. To obtain consumption levels note that we combine the first order conditions with the budget constraint. These are two equations in two unknowns, and therefore enable us to solve for C_1 and C_2.

Getting back to macro, savings in period 1, S_1, are:

$$S_1 = \begin{cases} Y_1 - \frac{Y_1+Y_2}{2} > 0 & \text{if lender} \\ Y_1 - \frac{Y_1+Y_2}{2} < 0 & \text{if borrower,} \end{cases}$$

so depending on your lifetime income stream, we could have $S_1 > 0$ or $S_1 < 0$. Of course, if (miraculously) $Y_1 = \frac{Y_1+Y_2}{2}$, there's no borrowing or lending in period 1 at all. And keep in mind that the supply of savings is the *demand for future consumption*; according to the PIH, savings are entirely for consumption-smoothing purposes.

Now let's talk about some implications. What's striking about the PIH is that *the timing of expected income is irrelevant.* If you expect to receive income next period, that should "kick in" today. So by the time you get the income, you will already have responded. Consumption-wise, *nothing* should change the day you get the cheque. Nothing. *There should be no reaction to anticipated income.* For this reason, consumption is often described as following a random walk. If a stochastic process follows a random walk, then all changes are unpredictable; there should be no expected changes. Another common example of a random walk are stock prices. For instance, if profits (and hence dividends) are expected to rise in the future, then stock prices should rise *today.* In particular, stock prices should not rise at the time dividends rise. Instead, they should rise and reflect this information beforehand. It is the same idea for the PIH.

Central to the analysis is the distinction between permanent and transitory changes. Returning to the example of the lotto prize, according to the PIH, how does consumption change upon getting the prize? If people base consumption on lifetime income, then a large prize this year should have little effect on consumption this year. However, to the extent it changes the present discounted value of all my income and hence *permanent income*, it *does* affect consumption to a small extent. That is, I will consume a little of the bonus this year, but will *smooth* the rest of it over my future lifetime. On the other hand, with a *permanent* doubling of income from, say, promotion, consumption rises permanently by the change. Thus, how much you consume is critically dependent on the *persistence* of the change in income. Can you see now the problem with the Keynesian function?

The PIH implies that you can infer a lot of information by watching the level of consumption in an economy. More broadly, one can view all of social insurance—such as pensions and "the dole"—are mechanisms to help people smooth their consumption. What's more, in international economics, countries behave like this too. A poor country—i.e., one with low consumption—might borrow or receive aid to finance development. Moreover, implicitly *you* are borrowing from the government. To pick a random example, the government finances your fees at Trinity, but you'll pay it back later in the form of higher tax payments. This way, the government helps you to smooth your consumption. See?

Assets

Suppose now the consumer starts off life with some level of assets, A (which could, say, be a bequest.) How does this change things? Because assets are a source of lifetime income, the intertemporal budget constraint is now

$$C_1 + C_2 = A + Y_1 + Y_2$$

Start with the budget constraint, and isolate C_2 to get:

$$C_2 = A + Y_1 + Y_2 - C_1.$$

Plugging this into lifetime utility yields:

$$u(C_1) + u(A + Y_2 + Y_1 - C_1)$$

Then maximize with respect to C_1:

$$u'(C_1) - u'(C_2) = 0 \Rightarrow u'(C_1) = u'(C_2) \Rightarrow C_1 = C_2$$

Again, we get the same result: consumption is the same in both periods. To find the *levels* substitute this back into the budget constraint to get

$$\boxed{C_1 = C_2 = \frac{A+Y_1+Y_2}{2}.}$$

Therefore, changes in the real value of assets affects consumption. Of course, A could also refer to the value of a consumer's portfolio or home. And in the last decade, changes in asset prices have had indeed had large impacts on consumption.

1.0.3 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:

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

The 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:

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

1.0.4 Liquidity Constraints

Practically all economists subscribe to some form of the PIH. Yet, in the data, consumption is moderately responsive to income changes; overall, smoothing is less than the PIH predicts. One common way to explain this within the framework of PIH is to invoke *liquidity constraints*. Point is, many people can't get loans. For instance, they're in debt already; they cant get collateral; they have criminal records, and so on.[4] And those who are *liquidity constrained* are stuck with the income they have. For this reason, with liquidity

[4] 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.

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 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.[5]

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 $\frac{1}{1-mpc}$.[6] 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—see e.g., Mankiw's blog—are essentially debates on whether the PIH is correct.

1.1.1 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.[7] *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, 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

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

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

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

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."

1.2 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.[8]

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:

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

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

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:

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

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

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.[9] 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.

1.2.1 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

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

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

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

[9]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 $C_=11$ and $C_2 = 13$ or when $C_1 = 9$ and $C_2 = 14$.

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:

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

1.2.2 Fisher Separation

Suppose now that the consumer can give up income, E, today and receive $w(E)$ tomorrow; e.g., a consumer can invest E in job training to receive a higher wage, $w(E)$ next period. Accompanying the fall in income in period 1 is a rise in period 2. Then the consumers problem is

$$\max_{C_1 \geq 0, C_2 \geq 0} u(C_1) + \beta u(C_2); \qquad \beta \in [0,1]. \tag{1.4}$$

subject to

$$\underbrace{C_1 + \frac{C_2}{1+r}}_{uses} = \underbrace{Y_1 - E + \frac{Y_2 + w(E)}{1+r}}_{sources}$$

The Lagrangian is

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

The first order conditions are

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

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

$$\frac{w'(E)}{1+r} = 1$$

The key insight here is that the consumer undertakes the efficient level of investment, *regardless* of the level of impatience, β.

1.2.3 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.

Finally, we can write the period budget constraint as

$$Y_t + (1+r_t)B_t = B_{t+1} + C_t$$

where $Y_t = w_t L_t + \pi_t$, where π_t are profits from firm ownership. Setting up the problem like this leads to a Lagrangian of the form:

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

We could easily solve for the optimal trajectory of consumption given initial assets B_0 and some transversality condition $B_{T+1} \geq 0$.

1.2.4 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.[10]

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 or yachts. The fact that the IES is low empirically suggests the average good is rather like salt.[11]

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

[10] 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.

[11] 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.

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?

Euler Equation with CRRA Utility

$$\frac{C_{t+1}}{C_t} = \left(\frac{1+r}{1+\rho}\right)^{\frac{1}{\theta}}$$

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 on the interaction of income and substitution effects. 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 since, this period's consumption may either rise or fall.[12] Happily for us, though, we *can* actually tell which way things go from the consumer's utility function; in particular, from—and you knew this was coming—the *intertemporal elasticity of substitution.* Because this tells us how extra 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 specif-

[12] 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.

ically 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 we'll have a *fall* in C_1 if r rises. And it's the other way round for $\theta > 1$; and of course effects just balance if $\theta = 1$.

Maximizing

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

subject to

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

gives

$$C_1 = \left(Y_1 + \frac{Y_2}{1+r} \right) \frac{1}{1 + (1+r)^{\frac{1}{\theta}-1} \beta^{\frac{1}{\theta}}} \tag{1.5}$$

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.)

From Eq (1.5), we see 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.

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.

1.3 Fiscal Policy and Ricardian Equivalance

So far, we have not internalized the government's budget constraint. In the previous analysis of fiscal policy, we found that the tax rebate would be smoothed over the consumer's lifetime and therefore the stimulus effect would be small. However, this was too optimistic. What happens if consumers know that *they* ultimately will have to pay for the tax break. Very briefly, the theory of Ricardian Equivalance (RE) answers this question: what happens to consumption if there's a tax cut today *and consumers know they will eventually have to pay for it?* To understand RE consider this: Imagine you are the only kid at TCD and, what's more, TCD's only income source is your fees. Furthermore, TCD has to power to raise fees at *any* time. Now suppose the college obtains a loan from the bank to give you a grant of 5000 euros. They must repay the loan next year. But nothing else changes, especially TCD's expenditure. Now ask yourself: how do you respond? Is that your final answer?

Well, what you should do is simply save the 5000 euros, since you know full well TCD has to repay the loan next year—and you are TCD's only income source! The same reasoning applies to a lump-sum tax-break from the government. There should be no income effect upon receipt of the lump sum tax break. In the context of government finance, this "neutrality result" is called the Ricardian Equivalence Theorem, or just simply Ricardian Equivalence (RE). To formalize this, suppose B denotes government borrowing, T lump-sum taxes and G government expenditure. As always, there is a single representative consumer with rational expectations. So anyway, in period 1 we have:

$$\underbrace{G_1}_{uses} = \underbrace{T_1 + B}_{sources}$$

and in period 2, noting the government have to pay back loans:

$$\underbrace{G_2}_{uses} = \underbrace{T_2 - (1+r)B}_{sources}$$

Combining these gives the government's intertemporal budget constraint:

$$G_1 + \frac{G_2}{1+r} = T_1 + \frac{T_2}{1+r} \tag{1.6}$$

Moving on, the consumer's problem is:

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

subject to:

$$C_1 + \frac{C_2}{1+r} + T_1 + \frac{T_2}{1+r} = Y_1 + \frac{Y_2}{1+r}$$

But the rational consumer knows that $G_1 + \frac{G_2}{1+r} = T_1 + \frac{T_2}{1+r}$; the consumer knows that he'll have to pay the bill for any government expenditure. So overall, the consumer solves:

$$\max_{C_1 \geq 0, C_2 \geq 0} u(C_1) + \beta u(C_2) \quad \text{subject to}$$

$$C_1 + \frac{C_2}{1+r} + G_1 + \frac{G_2}{1+r} + = Y_1 + \frac{Y_2}{1+r}$$

So the consumer's plans—in particular, consumption—are *completely independent* of the timing of taxation. All that matters is that presented discounted level of government expenditure. Really, this is just an application of the PIH. According to the PIH, all that matters for the consumer is the *present discounted value of lifetime income.* And—to come to my main point—this is unaffected by whether taxes come today or tomorrow. Hence, assuming government expenditure plans don't change, *the consumer spots the gimmick and* *will not respond to a lump-sum tax-break today.*[13] In other words, there's no income effect arising from the tax cut; the consumer does not any feel richer. Turns out, the tax-cut will be saved for the future tax hike. *End of story.* There's no avoiding the tyranny of equation Eq. (1.6).

Another way to see this is as follows. Say your taxes change today by $-T$. From the government's budget constraint, we have $G_1 + \frac{G_2}{1+r} = T_1 + \frac{T_2}{1+r}$, and we assumed there is no change on the left-hand side. So if T_1 falls by 10, say, then for equality to hold, T_2, must rise by $10(1+r)$. More generally, if taxes today fall by T, then tomorrow you will have to pay back $(1+r)T$. So, as a result of a tax break of T, the present discounted value of the change in your tax bill today—what matters for consumption—is:

$$\underbrace{-T}_{\text{value of tax-break today}} + \underbrace{\frac{T(1+r)}{1+r}}_{\text{PDV of future tax hike}} = 0$$

Hence, there's no change in lifetime income and consumption. The multiplier is zero.

Neutrality Theorems

Remember the Modigliani-Miller (MM) Theorem? In case you forget, the MM theorem says that, given certain assumptions such as perfect capital markets, the value of a firm is independent of whether it's financed by debt or equity. All that matters for firm value are the future cash flows of the firm, and these depend on fundamentals— things like real assets, growth opportunities, product development, and so on. Now here's the thing: RE is really the MM Theorem for government finance. Here's why. Ricardian equivalance says that whether the *government* is financed by *debt* or *taxation* doesn't matter from a consumer's standpoint. At the end of the day, the bill has to be financed by taxes (that would be *you*) *in any case.* Face it, it's tax now or tax later; so pick you poison, there's no free lunch. While all that matters for firm value is the expected path of cash-flow, *all that matters for the value of the tax-burden are the fundamentals; i.e., the path of government expenditure.*[14]

[13]Note that so far we've been talking about lump-sum tax rebates and <u>not</u> tax rates.

[14]An important implication of the MM theorem is that dividend policy is irrelevant, since the PDV of dividends must equal the PDV of cash-flow. Whether you get dividends today or tomorrow doesn't matter; only cash-flow does. Symmetrically with RE, the PDV of taxes equals the PDV of government expenditure. So whether taxes come today or tomorrow doesnt matter; only government expenditure does.