EC4010 Macroeconomics

January 2, 2013

Chapter 1

Consumption Theory

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*:

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

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

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

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

  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, let's start with 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, a class of students. 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:

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

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

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.[4] 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.*

Result 2 *According to the permanent income hypothesis, consumers smooth marginal utility over time. The concept of marginal utility is central to the decision to save and consume.*

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

[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), while the lifecycle hypothesis focusses on saving for old age.

[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.

[4]As a general point, if you income is relatively high, it could be mainly due to transitory income. This has implications for tax policy: namely, one's tax rate should depend on one's permanent income, not their transitory income.

the sandwich, it's reasonable to assume people act *as if* they're solving a maximization problem.

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.

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 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.[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. Since the rule holds for any two arbitrary periods, it holds for arbitrarily many periods too.

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.[6]

Below, I'll derive the optimal conditions with the Lagrangian technique.

[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.

[6] 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.

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:

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

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

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.[7] 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 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:

[7] 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 $C_2 = 14$.

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

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.[8]

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.[9]

[8] 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.

[9] 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

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?

essential—making it like salt, so to speak.

Chapter 2

Labour Supply

Now we turn to labour supply. Introducing labour supply into the standard two-period model means there are now effectively two goods: consumption and leisure. Here, leisure is effectively a consumption good that the household "consumes." By supplying labour and taking less leisure, the household is implicitly exchanging the leisure good for more regular consumption goods. The *period utility function*—i.e., the utility function at a given point in time—now takes the form:

$$u(c, l) = u(c) - v(l),$$

where c denotes consumption and l denotes labour supply. For the usual reasons, consumption exhibits diminishing marginal utility. Meanwhile, the distutility of supplying labor is convex; like climbing a stairs, it gets harder as the level rises. This way, the consumer will want to spread labour supply over time. In this environment, therefore, consumers desire to spread consumption over time, but also want to spread the labour over time. This eagerness to spread labor supply over time depends on how convex marginal disutility of labour is. For instance, the fact people take long vacations in summer and work five day weeks suggests labor disutility is not *that* convex.

Most importantly, people use the fruits of their labour to purchase consumption. You can think of labour as a means of attaining consumption. In this model, that's why people work. The fundamental tradeoff here is between labour, which provides disutility, and consumption which provides utility. For this reason, one key determinant of labour supply is how quickly diminishing marginal utility sets in. To see this tradeoff, suppose the budget constraint is:

$$c = wl + d,$$

where w is the real wage and d is other income, say dividends. For each period, consumer now solves:

$$\max_{c,l} u(c) - v(l) \qquad \text{subject to} \qquad c = wl + d.$$

Noting the dependence of c on w, I can maximize this using the chain rule. This gives the first order condition:

$$u'(c)\frac{dc}{dl} - v'(l) = 0 \Rightarrow u'(c)w = v'(l)$$

This is static neoclassical labor/leisure optimality condition, and is generally written as:

$$wu'(c) = v'(l), \tag{2.1}$$

that is, the utility gain to supplying an extra unit of labor is equal to the marginal disutility of labor. The real return to an extra unit of labour is just the wage multiplied by marginal utility. Namely, the real wage indicates how many goods I can purchase by supplying an extra unit of labour. And multiplying this by marginal utility, $u'(c)$, gives the total utility gain. Conveniently, we can think of this as a "marginal gain equals marginal cost" condition.

2.0.2 Income and Substitution Effects

How does a rise in the wage rate affect labour supply? Whether labour supply rises or falls (relative to the previous optimal plan) depends on the interaction of income and substitution effects. First, there is the *substitution effect*; as with all substitution effects, it deals with the change in relative prices. Now that the return to working is higher, you should work more. Put another way, a rise in the wage rate makes today's leisure relatively more costly. And this makes you consume *less* leisure today. In short, the substitution effect says: *go for it, work more*.

Second, there is the *income effect*. You are now earning more on all your *existing* hours of labour, and therefore you have more income performing the *same* amount of work. And seeing you are now richer, there's less need to work; you should consume more *leisure* today. So the income effect says: *look, you're now better off; work less.* Note that this effect only holds for those who are *already working*. For those not working, there is no income effect and only a *pure substitution effect.* For this reason, there is a lot more action on this extensive margin in response to changes in wages rates. For example, the rising wages of females can explain the rise in female labour supply over the past forty years. In addition, changes in tax rates have the largest effects on movements into the labour force.

Whether the income or substitution effect dominates depends on the form of the utility function. Looking at the optimality condition in (2.1) above, we can see the tension between the two effects. On one hand, the substitution effect is given by the rise in w. All else constant, a rise in w raises $v'(l)$; that is, a rise in w increases marginal disutility. Why? Because labour is rising; this is what the substitution effect dictates. On the other hand, the expression for marginal utility, $u'(c)$ mediates the income effect. All else constant, a rise in the wage causes c to rise. And given that this causes marginal utility to *fall*, the income effect dictates that $v'(l)$ should *fall*; i.e., labour supply should fall. How quickly marginal utility falls is therefore important for determining whether the income or substitution effect dominates. Long-run time-series and cross-sectional evidence suggests a dominant income effect. See Figures 2.1-2.3.

2.0.3 An Example

To give an example, suppose $u(C) = \frac{C^{1-\theta}}{1-\theta}$, $v(l) = \frac{1}{5}l^2$, and $c = wl$. As we know, the parameter θ mediates how quickly diminishing marginal utility sets in. Following through with this reasoning, θ mediates how quickly people are satiated and how strong income effects are. Now, the consumer solves:

$$\max_{l} \frac{(wl)^{1-\theta}}{1-\theta} - \frac{1}{5}l^2$$

The first-order condition is:

$$\frac{w}{(wl)^{\theta}} = l \Rightarrow l^* = w^{\frac{1-\theta}{1+\theta}}$$

Hence, if $\theta > 1$, $\frac{dl}{dw} < 0$, and the income effect dominates. Intuitively, since θ is relatively high, diminishing marginal utility sets in quickly. So why bother working for such low utility gains? What's the point? By contrast, If $\theta < 1$, the substitution effect dominates and labour supply will rise as the wage increases; that is, $\frac{dl}{dw} > 0$. Now, to liven things up, suppose we have a proportional tax rate, t, on labour. The optimality condition is now:

$$\frac{(1-t)w}{((1-t)wl)^{\theta}} = l \Rightarrow l^* = ((1-t)w)^{\frac{1-\theta}{1+\theta}}$$

In the case where $\theta > 1$, a high tax rate *raises* labour supply. By contrast, if $\theta < 1$, a high tax rate *reduces* labour supply. Instead of working more as tax rates rise, people now work *less*.[1]

Dynamic Model

In a two period dynamic model, where income is now *endogenous*, this consumer's problem is:

$$\max_{C_1, C_2, l_1, l_1} u(C_1) - v(l_1) + \beta\left(u(C_2) - v(l_2)\right)$$

To keep things simple, assume that the wage rate is constant over time. Denoting savings by S, the income constraint in period 1 is

$$\underbrace{wl_1}_{sources} = \underbrace{S + C_1}_{uses}$$

And in period two is:

$$\underbrace{wl_2 + (1+r)S}_{sources} = \underbrace{C_2}_{uses}$$

Combining these gives the *intertemporal budget constraint*:

$$C_1 + \frac{C_2}{1+r} = wl_1 + \frac{wl_2}{1+r}$$

Then the Lagrangian is:

$$\mathbb{L} = u(C_1) - v(l_1) + \beta\left(u(C_2) - v(l_2)\right) + \lambda\left(wl_1 + \frac{wl_2}{1+r} - C_1 + \frac{C_2}{1+r}\right)$$

Differentiating with respect to l_1, l_2, C_1, and C_2, and tidying things up, gives the optimality conditions:

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

$$wu'(C_1) = v'(l_1)$$

[1] An interesting policy question is, what tax rate maximizes the government's tax revenue? To address this question, note that tax revenue in this world is given by $twl^* = tw\left((1-t)w\right)^{\frac{1-\theta}{1+\theta}}$.

$$wu'(C_2) = v'(l_2)$$

Taxation

With *proportional* labour taxes of t on labour each period, the intertemporal budget constraint becomes:

$$C_1 + \frac{C_2}{1+r} = (1-t)wl_1 + \frac{(1-t)wl_2}{1+r}$$

Hence, the optimality conditions become:

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

$$(1-t)wu'(C_1) = v'(l_1)$$

$$(1-t)wu'(C_2) = v'(l_2)$$

The latter conditions are effectively instructions dictating the consumer's optimal labour supply. Because the prices faced by the consumer are now different, the optimal labour/leisure choice is *distorted* and different from before. Of course, the response depends on the interaction of income and substitution effects.

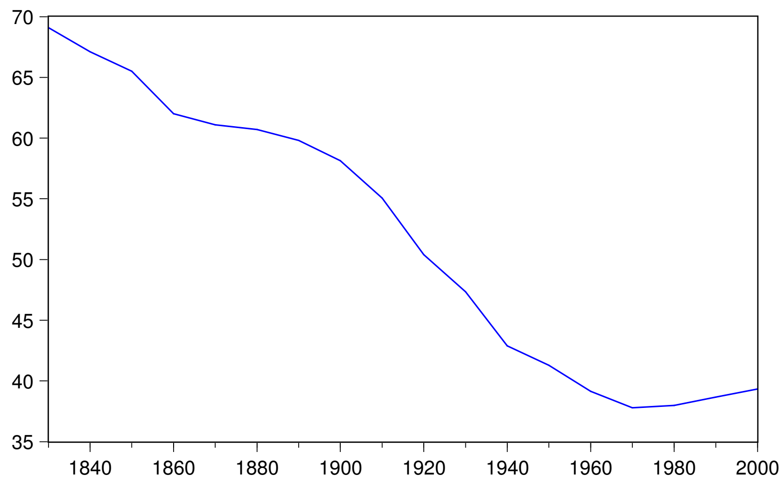

Figure 2.1: Weekly Labour Hours: U.S., 1830-2000

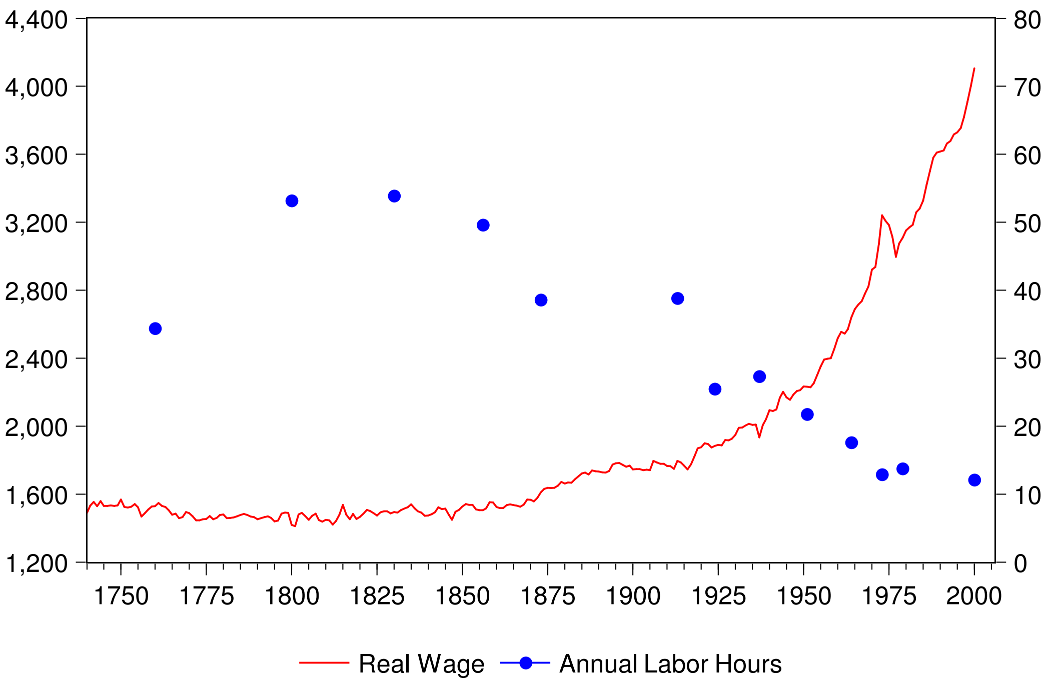

Figure 2.2: Annual Labour Hours and Real Wage Index: Britain, 1750-2000

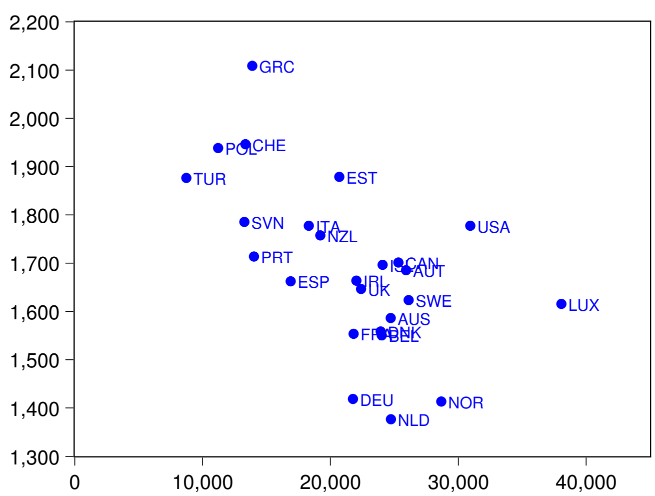

Figure 2.3: Labour Hours per Worker and GDP per Capita: OECD Countries in 2010
Source: OECD iLibrary

Chapter 3

Long-Run General Equilibrium Model

Up until now, we have taken the interest rate as given in our standard two period model. Yet in reality, nothing is simply *given* at the aggregate level: everything is endogenous. In a partial equilibrium analysis, we take prices—here, the interest rate—as exogenous. While individual consumers surely take the interest rate as given, at the economy-wide level the interest rate is a market-determined price. And like all prices, this is determined by the interaction of supply and demand. In a general equilibrium model, we solve for demands *and* prices. To start with, I want to analyze what is called the *natural rate of interest*. You can consider this the fundamental or average real rate that should prevail in an economy *on average*. For instance, if you were asked what the real rate will be in 100 years time, the answer would be the natural rate. It's the average rate the prevails when the economy is neither in boom nor recession. Therefore, implicit in what follows is the assumption that the economy is at potential, with output and demand given by *potential output* (a level determined by a long-run growth model such as the Solow model.) For simplicity, I assume there's no uncertainty, no growth, and no risk. There are free, frictionless markets, prices are flexible, and the classical dichotomy holds. And because I am not yet introducing money, think of everything in terms of goods.[1] For example, if I lend you 100 coconuts and you give me 104 back, then $r = 4\%$. Because of this long-run setting, please forget about the Federal Reserve, money, Keynes, recessions, IS-LM, booms and all that short-run material for now. In addition, I will ignore the role on uncertainty.

[1]Equivalently you could say that everything is in nominal terms, but prices are normalized to one. We simply get rid of prices, since money is neutral in this model.

We seek to find the natural rate at any point in time. For the determination of the natural rate, the supply of loanable funds comes from savings, while the demand stems from investment. As a result, the natural rate stems from the interaction of consumption and production decisions—and will ultimately depend of preferences and technology. In the model below, there is a representative agent and firm. I assume households place all of their savings in the financial markets. Investors then borrow these for investment—and the natural rate ensures equilibrium. As such, we are assuming financial intermediation works well: financial markets ensure the stock of savings flows into investment. Rather that saying savings and investment, it's common to talk in terms of bond demand and bond supply. That is, people save by demanding bonds, while firms invest by supplying or issuing bonds. The economy lasts for two periods. But, more generally, we can extend the analysis to arbitrarily many periods. As we already know, the standard Euler equation gives the optimal allocation of consumption across two periods, but this holds for *any* two consecutive periods for *any* arbitrary length of time. So the standard first order conditions pin down the optimal consumption path for any length of time; implicitly this gives us the savings decision. Restricting to two periods just simplifies the analysis. Anyway, lifetime utility is

$$u(C_1) + \beta u(C_2) \qquad u' > 0, \qquad u'' < 0.$$

The budget constraint in the first period is

$$C_1 + B_1 = Y_1.$$

Here, B_1 are bonds. We could also write this as savings and simply assume agents save through purchasing bonds. If B_1 is positive, the consumer demands bonds. This just means the consumer is saving, and $S_1 = B_1$.[2] In the second period, the budget constraint is

$$C_2 = Y_2 + (1 + r)B_1,$$

that is, the consumer consumes period 2 income plus the return on the bonds. Manipulating these constraints gives the consolidated or *intertemporal* budget constraint. As shown in Chapter 1, the standard lifetime income constraint is

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

[2]In this analysis, supply of savings equals demand for bonds.

So where does the income come from? Well, one obvious candidate is labour supply. But because we are not interested in production here, we won't model the sources of income, Y_1 and Y_2. Both refer to potential output. The consumer takes r as given and maximizes lifetime utility subject to the intertemporal budget constraint. The standard Euler equation is

$$\underbrace{u'(C_1)}_{pain} = \underbrace{\beta(1+r_1)u'(C_2)}_{gain}$$

Now, this is the usual *pain versus gain* condition.[3] At equilibrium the utility loss from giving up a little consumption equals the utility gain. I give up one good and lose $u'(C_1)$. Next period I get $(1+r_1)u'(C_2)$ goods. And since I discount the future by β, the utility gain to saving in period one is $\beta(1+r_2)u'(C_3)$. If this doesn't hold, we can shift around consumption a bit and do better. Implicitly, this pins down the optimal evolution of consumption. (Always bear in mind that *all* first order conditions pin down optimal demands (or demand curves).)

Combining the intertemporal budget constraint with the Euler equations give the optimal *levels* of C_1 and C_2, and in turn the optimal degree of saving in period 1. Yet, from a theoretical standpoint, it's not clear whether savings are increasing or decreasing in the real interest rate. Yet I assume $S'(r) > 0$; that is, savings are increasing in the real rate of return. Empirically, for any given saver, although there is a positive effect, the effect is small; i.e., $s'(r) \approx 0$.

The sensitivity of savings to the change in interest rates depends on the concavity of the utility function and, specifically, the *intertemporal elasticity of substitution*. This parameter determines how much the consumer is willing to shift consumption across periods and depends crucially on the degree of diminishing marginal utility to consumption each period; i.e., the concavity of the utility function. (To see what I'm talking about, content yourself that the reason you don't skip lunch today and have two tomorrow instead is that there's sharp diminishing marginal utility to lunches (making utility very concave). Therefore, your response to a large interest rate—say, give me your lunch today, and Ill give you five lunches tomorrow—is likely small.)

The Representative Firm

Turning now to the representative firm, the firm represents the *borrower* and supply of bonds. The firm is perfectly competitive—a standard assumption in long-run models. For this reason, the firm takes the price of its goods as given (that's why I simply normalize the

[3] If we had more periods, we'd have $u'(C_2) = \beta(1+r_2)u'(C_3)$ and so on.

price to one.) Again, everything is in real terms. The firm's profits derive from the number of items produced—given by production—less the cost of production–the purchase price of capital. Investment takes place in period 1, while production occurs in period 2.[4] There is no labour in the model. The firm starts off with no capital, but in period 1, the firm borrows an amount K (i.e., the investment), and promises to pay back $(1+r)K$ at the very end of the period 2. The firm purchases investment in period 1, so as to have capital in period 2. For this reason, K in period 2 derives solely from investment, I, in period 1. The firm takes r as given and maximizes the present discounted value of profits that will arise in period 2:

$$\pi = \frac{Af(K) - (1+r)K}{1+r}, \qquad f'' < 0.$$

This is a static problem: there are no dynamics to the firm's choice of investment.

Maximizing with respect to K gives

$$\underbrace{Af'(K)}_{gain} = \underbrace{1+r}_{pain}$$

Note that the above is an equilibrium condition, not a definition. Think of it as a *rule* that dictates the firm's optimal plans. Because $f'' < 0$, investment is decreasing in the interest rate. If the right hand side is high, the left hand side is also high; that is, the marginal product of capital is high, meaning capital demand, and hence investment, is low. According to this condition, therefore, when the interest rate is high, investment is low.[5] This implicitly defines the negatively sloped demand curve for investment, and in turn, demand for loanable funds. Combining this downwardly sloping demand curve with an upward sloping supply curve, $S(r)$, gives the natural rate, r.

General Equilibrium

At the natural rate we have

$$Af'(K^*) = 1 + r^*$$

and

$$u'(C_1^*) = \beta(1+r^*)u'(C_2^*)$$

[4] I'm implicitly assuming income in period 1 is given exogenously by Y_1, but in more realistic settings, it could be determined by labour, or an initial stock of capital.

[5] Note that this relationship is not specific to this model; it also holds in more sophisticated model such as Tobin's Q model of investment.

In equilibrium, everything is endogenous: the price, r^*, and demands, K^*, C_1^* and C_2^*. Combining the optimality conditions gives the condition:

$$u'(C_1^*) = \beta A f'(K^*) u'(C_2^*)$$

$$\frac{u'(C_1^*)}{\beta u'(C_2^*)} = A f'(K^*) \qquad \Rightarrow \text{MRS=MRT}$$

$$\underbrace{\frac{u'(C_1^*)}{\beta u'(C_2^*)}}_{MRS} = \underbrace{A f'(K^*)}_{MRT}$$

Recall that this is a standard general equilibrium condition. Now this is an old friend from the micro part of EC3010. Optimality dictates that the rate at which you substitute consumption across periods equals the rate at which it's technologically possible. If you think about it, this makes sense. Imagine technology permits you to transfer one good today into 3 tomorrow: this is what the marginal rate of transformation, *MRT*, tells you; it's what "the world" or God permits you to do. The MRS tells you the rate at which *you* personally *desire* to substitute one unit today for another tomorrow.[6] If it were 2, for instance, Id be happy to give up 1 unit today for 2 tomorrow. But at these figures, the MRT offers a great deal. Technology permits you to give up one today for 3 tomorrow, but you'd be happy to do the exchange at a rate of 1 for 2. It follows that you'll continue making these exchanges until the MRT=MRS (by continuing to make the exchanges, the marginal product of capital *falls* next period, while the *MRS* rises; the "transfers" end when both are equal.)

Taxation

With a tax on capital income, the condition above becomes

$$\frac{u'(C_1^*)}{\beta u'(C_2^*)} = (1-t) A f'(K^*) \qquad \Rightarrow \text{MRS} < \text{MRT} = A f'(K^*)$$

In this case, the First Welfare Theorem won't hold, and the equilibrium is not pareto optimal. Namely, the consumer is now facing a distorted price and the economy rests at a point where $\frac{u'(C_1^*)}{\beta u'(C_2^*)} \neq A f'(K^*)$. From our reasoning above, this means it is possible to alter production/consumption and make the consumer better off. In this sense, the tax is *distortionary*. Note that a lump-sum tax would not have this effect. Although the consumer

[6] For instance, if $\frac{u'(C_1^*)}{\beta u'(C_2^*)} = 2$, then you value today's consumption twice as much as tomorrow's (marginal utility/"hunger" is higher today); in other words, you would be willing to give up one unit today for *two* tomorrow.

would be poorer, the choices would still be optimal in equilibrium (i.e., it would not be possible to make the consumer better off.)

3.0.4 Remarks

Note that the natural rate is *jointly* determined on the production and consumer side. As an example, suppose investment increases, causing the natural rate to rise. The subsequent rise in the natural rate depends on how responsive consumers are to rising interest rates, as we move up along the savings curve. In other words, the attendant rise in equilibrium savings and investment depends on the *intertemporal elasticity of substitution* (IES). As a result of the rise in investment—in the case of an elastic supply of savings—consumption will rise over time. In the extreme case where consumers are completely irresponsive to rising interest rates (i.e., $IES = 0$), the savings line would be vertical—and all the burden of adjustment would fall in the interest rate (i.e., on prices rather than quantities.) That is, consumers save more today and ultimately consume more tomorrow. To take a related example, imagine we have identical two closed economics (A and B), except that investment demand rises in A, but not in B. As a result, consumption growth will be higher in A (as people save more in response to higher interest rates induced by greater investment demand.) Fundamentally, the different consumption behaviours will be due to different investment rates in both countries. Another point to note is that different marginal products of capital across countries do not necessarily reflect technological advantage. To see this, suppose we have two different closed economies—A and B—where investment demand is the same in both, but savings are higher in A. In equilibrium, the marginal product of capital is higher in B. Yet this does not mean B has some technological advantage. Instead, it means that in equilibrium, the relative scarcity of savings in B means lower investment in equilibrium—and so the marginal product of capital will be higher there. Moreover, because parameters such as β lower savings, the high equilibrium marginal product of capital in B would be attributable *fundamentally* to a lower value of β in A. Point is, in a general equilibrium model, prices depend on structural features of the economy.

From the permanent income hypothesis we know that consumption—and hence savings—does not vary *that* much over time. For this reason, most changes in the natural rate stem from changes in the level of investment.

Output is always fixed at potential in this long-run model. All we are concerned about is the *distribution* of output among government expenditure, consumption, investment and net exports. The interest rate is the key to adjustment. For instance, if consumption increases, then saving falls and the interest rate rises. Investment then falls so we again have aggregate demand equal to potential again. Output is always *fixed* at potential. If one component of demand falls, a fall in interest rates will "crowd in" another part. Similarly, if one component of demand rises, this will "crowd out" another part.[7]

You might wonder, if the stock of savings increases over time as the economy grows, does this mean that the natural rate falls over time? No. As we know from long-run growth models, productivity is the main source of rising living standards. And according to these growth models, a rise in productivity leads to more savings *and* investment (more savings, since you are saving a roughly constant portion of a bigger pie; and more investment since new technologies are continually raising the marginal product of capital.) Empirically, therefore, savings and investment/GDP ratios are roughly constant over time, and the natural rate is fairly stable over time.

[7] Contrast this with short-run Keynesian models, where a rise in the savings rates induces a recession.

Chapter 4

Monetary Policy

4.1 Money Neutrality

When prices are flexible, a permanent doubling of the money supply just doubles prices in both periods. This is just the classical dichotomy. In this case, you can think of everyone receiving the money and running out to buy goods. But because output is fixed in the long-run, this increase in demand just causes prices to rise. *That* increase in prices then generates an increase in nominal (but not real) money demand to meet the new money supply. Nothing real changes. In particular, the natural real (and nominal) interest rate is still pinned down by the long run model, and output is pinned down by, say, the Solow model. With flexible prices, the money supply simply determines the price level: with flexible prices, a doubling of M induces doubling of P.

$$\frac{M}{P} = L(i, Y)$$

$$\frac{M}{P} = L(r_n + \pi, Y)$$

By contrast, if prices are sticky, then a permanent increase in the level of the money supply will cause i (and hence r, assuming long-run expectations of inflation are fixed) and Y to adjust. In this case, if there is excess supply of money, people try to unload it in the bond market. This puts downward pressure on interest rates. In turn, the corresponding fall in interest rates encourages people to demand money (since the opportunity cost of holding money has fallen).[1] In equilibrium, everyone must in fact hold the money, so we end up

[1] It's convenient to think of *you* getting a helicopter drop and then placing the money in the bond market. The attendant fall in interest rates then causes others to increase their money demand. Ultimately, money demand

with higher money demand and lower real and nominal interest rates. The attendant fall in interest rates is called the *liquidity effect*. We also get an increase in output when people try to unload the money on the goods market, *and* output is demand determined; this rise in output also raises money demand. But as we will see in the New Keynesian model—which gives a microfoundation to the above story—prices will eventually rise, causing interest rates to return again to their initial level. As a result, money is neutral again in the long-run.

4.2 Short-Run Interest Rates

Having shown how money can affect interest rates in theory, it's time to move beyond helicopter drops and talk about how the FED controls the money supply.[2] The FED controls the federal funds rate, i^*; this is its policy instrument. By buying securities from banks in *open market operations* and giving them dollars in return, it can increase the amount of reserves banking system (and vice versa). In turn, this increases the amount of reserves on the federal funds market, which lowers the federal funds rate. Then, if inflation expectations and prices are fixed in the short-run—the standard assumption—the real federal funds rate will also fall. It's important to keep in mind that the FED does not set rates. Rates are market determined, but the FED manipulates the market by changing the amount of reserves in the banking system. (Central to this story is that banks are compelled by law to hold reserves, but run out of them regularly, so need to borrow.) The FED has monopoly power over the creation of reserves—i.e., printing money and increasing the monetary base—which is key to its power. However, none of us ever pays the fed funds rate. When banks can borrow money/reserves cheaply, they typically pass that on to consumers. (After all, banks compete to make loans to customers; it's usually in their interest to attract borrowers.) Moreover, with access to cheaper reserves, they will lend more, increasing the money supply.

The fed funds rate moves closely with all short-run interest rates. Why? Well, the fed funds rate is the rate at which banks can lend or borrow reserves from each other. So suppose a bank can lend reserves to another bank overnight at 3 percent, the fed funds rate. But instead of lending to another bank, it can also buy some other short-term financial instrument—say a (hypothetical) two-night Treasury bill. For these fairly riskless markets to be in equilibrium, these rates must move closely. To see why, suppose we have a sit-

will increase until everyone is happy holding the new supply of money.

[2] Henceforth, I use the "FED" to represent *any* central bank.

uation where Treasury bills yield a return of 12%, while the fed funds rate is 5%. If this was the case, all banks would buy Treasury bills. This increase in demand for Treasuries would raise their price and lower their return towards the fed funds rate. Such arbitrage operations means that all short-term instruments of similar risk must pay approximately the same return in equilibrium. It follows, therefore, that a lower fed funds rate will tend to reduce almost all short-run rates.[3] For this reason, it's convenient and common to say the FED "controls" short rates.

4.2.1 Taylor Rule

The Taylor rule ensures the FED minimizes its loss function, which is increasing in the deviation of both inflation from target and output from potential. According to the Taylor Rule, the target for the fed funds rate is:

$$i^* = 2.5 + \pi + .5(\pi - \pi^*) + .5(y - y^*)$$

For example if inflation exceeds target the FED will try to contract the economy and lower the real rate below the natural rate and induce a recession.[4] Importantly, $\frac{\partial i^*}{\partial \pi} = 1.5$. So if inflation rises by 1 percent, the FED will raise rates by 1.5%. This way, it will raise real rates, $r^* = i^* - \pi$ by .5 in response to inflation. This idea of raising real rates when inflation rises is called the *Taylor Principle*: the increase in the nominal rate must be sufficiently high that it raises the real rate. In practice, most banks stress that inflation is their prime mandate and don't like to be perceived as targeting output. (Recall the dynamic inconsistency problem, whereby this can generate inflationary expectations.) Finally, to ensure stability in financial markets—and in particular bond prices—banks typically engage in *interest rate smoothing*. Of course the central bank faces a constraint, $i \geq 0$, that is sometimes binding (i.e., a liquidity trap).

Getting Real

Implicit in the Taylor rule is a target *real* rate:

$$r^* = 2.5 + .5(\pi - \pi^*) + .5(y - y^*)$$

Note that the FED targets the neutral real rate on average. Specifically, if output is at

[3] Via the term-structure equation, this in turn will affect all long-run rates. It is in this sense that the FED has leverage over the entire term structure of interest rates.

[4] The natural rate is the rate consistent with output at potential.

potential and inflation at target, the FED will aim for the natural rate. The figure 2.5 above is an estimate of the natural rate. Since the natural rate varies, however, a more general way to write this is

$$r^* = r_n + .5(\pi - \pi^*) + .5(y - y^*),$$

where r_n denotes the natural rate.

4.3 Long-Run Interest Rates and The Yield Curve

What really determines economic activity are long-run real rates. Yet the FED only controls the short-run rate. The term structure of interest rates shows how the FED can affect long rates too. The main theory of long-run rate determination is the *expectations theory of the term structure* or *expectations hypothesis*. From now on, this is the main theory we will use when we talk about long rates.[5]

Imagine you have to invest today for two years. If you have to invest now, there are two ways to get money to the same location: invest in a 2 year bond or buy a one year bond this year and then again in the following year. Assume the two year bond pays i_{2l} a year. The one year bond pays i_i this year and you expect it to pay i_2 next year. By an arbitrage argument, these two ways of investing a euro should earn the same return. Therefore,

$$(i + i_{2l})(1 + i_{2l}) = (1 + i_1)(1 + \mathbb{E}i_2)$$

$$(1 + i_{2l})^2 = (1 + i_1)(1 + \mathbb{E}i_2)$$

Taking logs

$$2\log(1 + i_{2l}) = \log(1 + i_1) + \log(1 + \mathbb{E}i_2)$$

For small x, we have the approximation $\log(1 + x) \approx x$. As a result,

$$2i_{2l} = i_1 + \mathbb{E}i_2$$

[5] Another theory is the *market segmentations hypothesis.* According to this theory, bonds of different maturity are not necessarily substitutable and, as such, the markets are segmented or independent. For example, big market players like pension funds or governments might demand long-run bonds, irrespective of their yields. In this case, the yields on long bonds would be artificially low and not necessarily reflective of expectations of future short rates; indeed, in this case, rolling over short run bonds would be more more profitable.

$$i_{2l} = \frac{i_1 + \mathbb{E}i_2}{2}$$

This holds more generally, so

$$i_{nl} = \frac{\sum_{z=1}^{z=n} \mathbb{E}i_z}{n}$$

According to this *expectations hypothesis*, the interest rate on a n year bond equals the average of the expected one-year returns over the next n years.

So far I've assumed risk neutrality; that is, investors only care about expected returns. This makes a long-run bond and rolling over two shorts perfect substitutes. Over long periods, inflation can be quite variable, making the fixed nominal payments associated with long-run bonds risky; in addition, prices of long-run bonds respond more to changes in interest rates. To account for risk, we must add a risk premium

$$i_{nl} = \frac{\sum_{z=1}^{z=n} \mathbb{E}i_z}{n} + \rho$$

Note that ρ typically depends on the maturity of the bond. To account for risk, we'd have $\rho \equiv \rho(n)$, where $\rho'(n) > 0$. More generally, the premium could compensate for liquidity differentials (most likely, it's easier to sell bonds of shorter maturity).

$$i_{nl} = \frac{\sum_{z=1}^{z=n} \mathbb{E}(r_t + \pi_t)}{n} + \rho.$$

The Yield Curve and Monetary Policy

The FED has considerable control over short rates, and hence more control over the short end of the yield curve. The FED can only affect long rates to the extent it affects short rates now and expectations of short rates over the next year or two. That's why communication, "open mouth operations" and, more generally, the management of expectations are all central to monetary policy. The expectations hypothesis is consistent with the fact that when the FED lowers short rates, long rates don't move as much. Relatedly, the hypothesis is consistent with the fact that short-run rates are more volatile than long-run rates (namely, long-run rates are an average and therefore are less volatile.) The risk premium makes the theory consistent with the typically upward sloping yield curve.

For practical purposes, I would say the yield curve reflects FED policy for 1-2 years out and expectations of the natural rate plus inflationary expectations thereafter. Why? Well, our best guess of the state of the economy after 2 years or so is potential output.

And when the economy is at potential, the interest rate equals the natural rate. In general, movements in long rates—with short rates held fixed—are typically reflective of changes in inflation expectations and risk premia. By adopting policies such as inflation targeting, central banks can control expectations of inflation right across the term structure. And to the extent such sound monetary policies can also reduce the inflation risk premium, they would also lower real long-term rates. I outline some example of yield curve movements below. Yet, keep in mind, that the yield curve is a perennial source of enigma to economists: many of its movements are inexplicable in terms of economic theory.

1. If FED increases money growth today, short-rates would fall. But long-rates could in fact *rise* if expectations of inflation rise too. (This represents a combination of the liquidity and Fisher effects.) This would happen if policy was deemed permanent. Conversely, if the FED sets short rates high now, long rates might come down in future due to lower risk premia and expectations (if this was a sign the FED is serious about fighting inflation).

2. Romer's text describes how high short rates today often lead to high long rates too—since markets often perceive unusually high rates by the FED as a signal that the FED has information about future inflation that they don't have. Hence as a precaution, they demand a higher risk premium on long-term bonds.[6]

3. Expectations of Clinton budget surpluses lowered long-run rates. Conversely, expectations of future deficits might raise long-run rates today.

4. Debt monetization would raise long-run rates due to expectations of inflation (and attendant increase in risk premium.)

5. In a currency crisis, the central bank typically raises short-run interest rates to attract inflows of capital. In this case, the yield curve would likely invert and slope downwards.

Bond Prices and Yields

An important relationship is the inverse relationship between bond prices and bond yields/interest rates. To see this (in the simplest way possible), suppose we have a one year bond: this year I pay P for the promise of $D > P$ next year. The key here is that D is fixed; hence a bond is

[6] In a similar example of strange asset price movements, sometimes lower interest rates by the FED cause the dollar to *appreciate*. Namely, investors take it as a signal that the FED is doing all it can to stabilize the economy—thus making it a more attractive place to invest.

often referred to as a *fixed income security.* The return on the bond (or its yield to maturity) is implicitly given by the relationship

$$P(1+r) = D$$

or equivalently, the price of the bond is given by

$$P = \frac{D}{1+r}$$

where r is the return on the bond. Now, if the FED raises interest rates from r to r', say, then the return on the bond must rise (by the usual arbitrage relationships). Therefore, *to generate* this higher return, r', the price of the bond will be

$$P'(1+r') = D \qquad P = \frac{D}{1+r'}$$

Content yourself the $P < P'$; that is, the price of the bond is now lower.

Another possibility is you buy the bond and two months later (say) the FED raises rates. According to the previous reasoning, the price of your bond now must fall; if you were selling it, it would certainly sell for less than P; namely, *to generate* the new higher required rate of return, you will have to lower the price of the bond. So if you intend selling the bond, this is clearly *bad news* for you: you will make a capital loss (formally, your *holding period return* would fall.)[7] It follows from all of this that bondholders pay close attention to interest rate movements.

4.3.1 Transmission Mechanisms of Monetary Policy

An important question is how monetary policy affects the real economy. Although the standard channel is through interest rates, there are a multiple of other ways too.

1. Interest rate channel. Lower interest rates lower the cost of capital. This increases investment demand. The purchase of (interest-sensitive) consumer durables also rises if rates fall.

2. By affecting risk-free rates, the FED can change a wide array of asset prices, since all of these are related to the risk-free rate. For example, recall the basic dividend discount formula: $P = \sum \frac{D_t}{(1+r_t+\rho)^t}$, where P is the stock price, r is riskless rate, ρ is risk premium, and D_t is the dividend in period t. The FED affects r_t, and a fall

[7] If you are holding the bond until maturity, this issue is of no relevance.

in risk-free rate therefore raises stock prices. Also, by preventing a recession, lower rates might increase profits and therefore expectations of D_t. In addition, the risk premium may fall. For all these reasons, stock prices rise.[8] Equivalently, think of any asset pricing model such as the CCAPM, which predict $r_i = r_f + \rho$, where r_i is the expected return on a stock and ρ is the model's predicted risk premium. Thus a fall in the risk-free rate, r, induces a fall in the stock's return, r_i. But this means the stock price must rise (recall that returns and prices are inversely related). This is why stock prices typically rise when FED cuts rates (in reality, prices will have risen in expectation beforehand, though, if the move was expected.) On this note, many say FED should raise rates to bring down stock/house price "bubbles." Yet this idea is highly controversial. Who is Bernanke to say a rise in stock prices represents a "bubble," rather than fundamentals? According to the *efficient markets hypothesis*, the stock price is an accurate measure of firm's worth. But this view might well change: the worst economic downturns of the twentieth century have all been associated with a collapse of asset price "bubbles."

Anyway, changes in asset prices have lots of indirect effects:

(a) Higher stock prices lower cost of capital since they make it cheaper for a firm to raise equity. Namely, by issuing shares, the firm now makes more revenue. In turn, this increases investment. This is called Tobin's Q theory. For the same reason, a rise in house prices stimulates housing production. Monetary policy raises house prices by making mortgages cheaper and thereby increasing demand.

(b) Higher stock/house prices raise the value of one's portfolio and lifetime wealth. By the permanent income hypothesis, this raises consumption.

(c) When asset prices rise, households have more liquid wealth, so they are "free" to buy more illiquid assets (e.g., consumer durables.) (Households would typically wish to hold liquid assets in case of, say, a medical emergency.)

(d) Balance Sheet effects (banks and lenders). The key here is adverse selection and moral hazard. Most U.S. firms rely on banks (and internal finance) rather than stock issuance. As a result, they are dependent on bank's willingness to lend. When interest rates fall, firm's profits typically rise (due to greater aggregate demand). Also, if firms are on adjustable loans, their debt payments fall, again raising profits. These increases in cash flow improve their balance sheets, so they can

[8] In addition, a fall in the interest rate would make bonds relatively unattractive, inducing people to purchase equities, thereby raising their prices.

offer more collateral to banks, raising banks willingness to lend. Also, greater collateral reduces the risk of moral hazard (i.e., the fear firms might use funds recklessly and default) since firms lose now lose more under default. Again, this makes banks more willing to lend. (This idea that banks make more loans in good economic times leads—leading to further good times—is what's called a "financial accelerator.") The same reasoning applies for households. For consumers on adjustable rate loans, lower interest rate costs increase reduce their short-term debt burden which again improves balance sheet.

3. The lending channel. Most importantly, this is different from the interest rate channel, since it represents a response on the *supply side*, not the demand side. With cheaper access to reserves in a monetary expansion, banks will be able and willing to lend more, which in turn leads to more investment. This view also emphasizes that banks will lend more if they have more capital themselves; for example, if banks suffer losses from default, they will lend less. According to this view, small firms are especially affected by banking problems, since they are reliant on bank lending (small firms rarely issue equity.) In addition, it stresses that banks have considerable "informational capital," (credit records of clients etc) which is essential for financial intermediation. Therefore, a collapse of banks destroys all of this and will reduce overall lending in the economy. Both the lending channel and balance sheet effects come under the heading of "the credit channel."

4. Lower interest rates reduce *credit rationing*. Recall that high interest rates increases the adverse selection problem since they attract riskier clients. Low interest rates alleviated this problem and make banks more willing to lend.

5. When the interest rate falls, the exchange rate depreciates. This stimulates exports and improves the current account.

6. A greater money supply will ultimately raise the price level and in turn reduces the real value of nominal debt. This can improve debtors balance sheets, making them better lending propositions.

4.4 Issues in Monetary Policy

- Liquidity Trap: this is when the fed funds rate hits zero. One way out is foreign exchange intervention; i.e., print money and use to purchase foreign currency. The associated increase in domestic currency causes a currency depreciation, which should improve the current account. The bank can also buy long-run bonds to reduce their returns and lower cost of borrowing. More generally, the FED can buy a range of assets so as to increase the money supply. Central bank can generate expectations of large money growth/inflation in the future to increase demand today. (A corollary of this is for the government to run large deficits to generate expectations of debt monetization.) Because $r = i - \pi^e$, such expectations of inflation can lower real rates for a given nominal rate and thereby stimulate economy. Paradoxically, Japan couldn't do this, since they had (too) successfully generated expectations of low inflation to combat dynamic inconsistency problem. Finally, one extreme solution is to make money bills go "obsolete" after a year (say) to encourage spending today.

- Why to banks target positive inflation?

    1. First, targeting non-zero inflation gives monetary policy more leverage. Suppose natural real rate is 1%. If inflation was 5%, we know the short-run nominal *natural* rate will be 6%. By contrast, if inflation was 0%, the short-run nominal natural rate would be 1%. Because the FED targets the nominal natural rate *on average*, it'd have greater "power" to reduce rates if inflation were higher. In particular, it can lower real rates. With low average rates of inflation, the risk of falling into a liquidity trap rises.
    2. Second, inflation can reduce real wages if nominal wages are sticky. (this is called "greasing the wheels" of the labour market to make it more flexible.)
    3. Third, a positive rate of inflation also makes deflation less likely. Deflation is dangerous since it encourages households to put off spending to future, thereby reducing spending today (leading to more deflation and creating a vicious circle). Central banks try to avoid deflation at all costs. Deflation typically associated with "depression-like" economies like Japan's "lost decade" in the 90s and the Great Depression.
    4. High inflation can be used as a source of finance for the government. Because raising tax rates on capital or labour might be highly distortionary (costs of taxation on anything are convex), this can be an efficient source of revenue.

- Preemptive policy: banks act early and react to inflation *forecasts*. If inflation rises, it generates expectations of inflation, which can be hard to counter.

- In the famous words of Milton Friedman, money affects output with "long and variable lags." Monetary policy is widely believed to be effective than fiscal policy.

- Money Multiplier: So far, I've assumed the money multiplier is constant and increases in the monetary base lead to increases in the money supply. This typically happens but sometimes does not. In particular, there's no guarantee banks will actually *lend*. As such, the Fed only controls the money supply *indirectly* via the monetary base. What economists mean by the money supply is the total money created from the initial monetary base. With banks lending money and consumers depositing it etc, the monetary base is continually "recycled." The money supply is given by $M = \mu mb$, where μ is money multiplier and mb is monetary base. The monetary authority controls the monetary base; and cet par the money supply; this normally works fine. For a high multiplier, we need banks holding few excess reserves and households depositing their money at banks (not under mattress!) Falls in the money multiplier are invariably associated with severe downturns (Japan 90s, Great Depression). See the credit channel discussion above for why the multiplier is procyclical and for why the money supply rises endogenously in booms.

- Why is monetary policy less effective in inflationary environment? Prices more flexible since people are more "switched on" to policy regime. Example of Lucas Critique: response coefficients change in different policy environments. Consequence: monetary policy should be more powerful when people don't expect it (i.e., in low inflation environments.)

Chapter 5

The New Keynesian Model

5.1 The Model

The New Keynesian model is the benchmark model used by central banks and most economists today. Ironically, it was spawned in response to RBC theorists, who claimed—with some justification—that old traditional Keynesian models—like the ISLM or Keynesian cross—were too ad-hoc to be taken seriously and in particular, didn't stem from fundamental microeconomic relationships. The key features of the model are

1. Consumer optimization, where labour and consumption decisions are results of maximizing lifetime utility.

2. Firm optimization, which leads to labour demand and optimal pricing/production. For simplicity, there is no capital. There is an imperfectly competitive product market and firms choose prices. However, they take wages as given; these are determined in the labour market, where wages are flexible and the market clears. Because of this, there is no unemployment in this model. (Numerous extensions exist to modify this unrealistic feature.)

3. Market Clearing/Goods Market Equilibrium. The model is a general equilibrium one, in that all markets will clear.

4. There is a representative household, N monopolistically competitive firms, and a government. The role of the government is unimportant here.

5. Demand-Determined Output (up to a point anyway): production ultimately responds to demand.

6. Because price-stickiness is central to money nonneutrality, it plays an important role in the model. There is time dependent pricing (not *state*-dependent, which would make prices dependent on the state of the economy.) At any given point in time, some prices will be fixed (regardless of events): firms irregularly set prices. We will rationalize this via "menu costs."

7. To have pricing decisions, we need price setters. For this reason, the market structure is *monopolistic competition.* This is in stark contrast to the price-taking assumption under perfect competition, which is more or less standard in modeling the long-run.

8. The fact $p > mc$ is key. This means there is the potential for demand-determined output. Because firms' prices are initially higher than marginal cost, they will find it optimal to increase production in response to an increase in demand, assuming prices remain fixed. There is certainly leeway to do this.

9. We need to get back to long-run level of output/potential. The *natural rate hypothesis* holds. The Phillips curve relationship will bring us from the short run to the long run.

The model has three key equations

- The *New Keynesian IS Curve* which derives principally from the equilibrium Euler equation plus some exogenous sources of demand (say, exports.)

- The *Taylor Rule*, which determines the interest rate. This does away with LM curve. Here, the money supply is endogenous; the FED just adjusts the money supply to hit the rate dictated by the Taylor rule. (You could think of the money market equilibrium condition, $M^s = M^d = L(i, y)$, in the background, where the FED is adjusting M^s to hit its target i.) Importantly, we ignore issues relating to falls in the money multiplier etc. Many version of the model incorporate the idea of *interest rate smoothing*; i.e., the rate chosen by the monetary authority is a weighted average of the existing rate and what the Taylor rule dictates.

- *The New Keynesian Phillips curve.* This will describe the adjustment of prices (and output) until the economy reverts to the natural/potential rate of output again.

5.2 The Household

The household maximizes:

$$E_0 \sum_{t=0}^{t=\infty} \beta^t \left(\frac{C_t^{1-\theta}}{1-\theta} - \frac{L_t^{1+\sigma}}{1+\sigma} \right)$$

The important parameters here are θ, which governs diminishing marginal utility of consumption and σ, which controls the marginal disutility of labour; as such this parameter mediates the household's incentive to smooth labour supply over time. Think of the σ as the "oh my back hurts" parameter: it mediates how painful further hours of work are. The parameter, $\beta = \frac{1}{1+\rho}$, is the consumer's discount factor. Related to this is the parameter ρ, which is the consumer's *rate of time preference*. Now, the marginal utility of consumption is

$$u'(C_t) = \frac{1}{C_t^{\theta}}$$

The marginal disutility of labour is

$$L_t^{\sigma}$$

Of course, if $\sigma = 0$, people would be *indifferent* to working, say, 100 hours one day and 50 hours a day for two days.

The flow budget constraint at time t—there is one every period—is

$$\underbrace{W_t L_t + (1+i)B_{t-1} + \Pi_t}_{sources} = \underbrace{P_t C_t + B_t + T}_{uses}.$$

Π_t denotes profits from the firms at time t. We assume the household owns the firms (after all, *someone* owns them, and the profits have to go somewhere). B_t refers to bonds, which I assume are issued by the government to finance its expenditure; i refers to the *nominal* interest rate. At the start of time the number of bonds, B_0, is given. We could aggregate all the flow budget constraints into a intertemporal budget constraint, like we did in the two-period model. As it is, though, we can maximize the objective function subject to all the flow budget constraints using the technique of Lagrangian multipliers. We can use the Lagrangian technique to solve this, bearing in mind that there is a constraint for each time period. Fortunately, it is relatively easy to solve. To solve it, I'm just going to pick just two random periods t and $t+1$ and invoke the familiar conditions from our two-period model. Restricting the analysis to two periods like this is without loss of generality. However, when dealing with infinite time, one also must impose a *transversality* condition to ensure the consumer doesn't continually permit B to grow indefinitely large or negative. This way, the consumer is not permitted to die in debt; in addition it also ensures the consumer consumes everything in the "last period." Leaving positive assets at the end would obviously not maximize utility if the consumer derives utility from consumption; as such, this also acts as an optimality condition.

The form of the utility function is

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

The parameter θ mediates the degree of diminishing marginal utility and hence the consumer's willingness to shift consumption intertemporally. In particular, the parameter, $\frac{1}{\theta}$ is the *intertemporal elasticity of substitution* and is a measure of the consumer's willingness to shift consumption across periods. Because of this, the parameter will determine the households response to interest rate changes. For example, if $\frac{1}{\theta}$ is high—i.e., θ is low—then the consumer will be eager to save in response to increases in interest rates. From now on, I will assume that this is case and that the substitution effect of price changes dominate the income effect. (If you think about it, this makes sense for a business cycle model: the prices changes are typically temporary, making the attendant income effects weak and substitution effects strong.)

Optimality Conditions

The first order condition for consumption in period t is

$$\underbrace{\frac{1}{P_t}u'(C_t)}_{pain} = \underbrace{\beta\frac{1+i}{P_{t+1}}u'(C_{t+1})}_{gain}$$

Think about it: You give up a euro today. Hence you give up $\frac{1}{P_t}$ *goods*. Because the value of those goods was $u'(C_t)$ each, you therefore lose $\frac{1}{P_t}u'(C_t)$ in utility. Next period, you gain $1+i$ back (initial sum plus interest). With that you can buy $\frac{1+i}{P_{t+1}}$ goods, giving you extra utility of $\beta\frac{1+i}{P_{t+1}}u'(C_{t+1})$. You get $u'(C_t)$ in each unit, and since we value the future less, we discount everything with β.

Because the consumer faces uncertainty about the future (the future price level or future consumption) we should put an expectation operator the right-hand side, giving:

$$\frac{1}{P_t}u'(C_t) = \mathbb{E}_t\beta\frac{1+i}{P_{t+1}}u'(C_{t+1})$$

From now on, however, I'll omit this. Because uncertainty doesn't play a significant role in the analysis, this is fine. Tidying up the above and omitting the awkward expectations operator, we get

$$u'(C_t) = \beta \frac{P_t(1+i)}{P_{t+1}} u'(C_{t+1})$$

Because $\frac{P_{t+1}}{P_t} = 1 + \pi_t$, this gives

$$u'(C_t) = \beta \frac{(1+i)}{1+\pi_t} u'(C_{t+1}).$$

Noting that $\frac{1+i_t}{1+\pi_t} \approx 1 + r_t$,[1]

$$u'(C_t) = \beta(1+r_t)u'(C_{t+1}),$$

where $1 + r_t = \frac{1+i_t}{1+\pi_t}$ is the gross real rate of interest. This governs the *path* of consumption. There is an Euler equation each period. (To get the actual *level* of consumption at a point in time, we'd have to combine the Euler equations with the budget constraints. But the steepness of the consumption (i.e., the rate of consumption growth) is the same for everyone. So if you're a millionaire say, then the steepness of your consumption profile will be the same as a poor person; but of course the levels of consumption in each period *will* be different.) Moving on, the labour/leisure optimality condition is

$$\frac{W_t}{P_t} u'(C_t) = v'(L_t) \Rightarrow \frac{W_t}{P_t} \frac{1}{C_t^\theta} = L_t^\sigma$$

The *transversality condition* is

$$\lim_{t \to \infty} \beta^t u'(C_t) B_t = 0$$

Basically, this means you don't want to leave any savings (i.e., bonds) left at the end if you value consumption. Of course, if you don't value consumption "at the end," then $u'(C_t) = 0$, and then it's ok to leave positive savings left over (the TVC allows this).

5.2.1 Demand

Note that, given the functional form for utility, $u'(C) = \frac{1}{C_t^\theta}$. Marginal utility is falling in the level of consumption, and the extent to which it falls depends on that important parameter, θ. Substituting this into the Euler equation above gives

[1]To see this formally, take logs of both sides to get $\log 1 + r = \log(1+i) - \log(1+\pi_t)$. Then noting the approximation $\log(1+x) \approx x$ confirms that $r = i - \pi$.

$$C_t^{-\theta} = \beta(1+r_t)C_{t+1}^{-\theta}$$

Taking logs gives

$$-\theta \log C_t = \log \beta + \log(1+r_t) - \theta \log(C_{t+1})$$

Letting $c_t = \log C_t$

$$c_t = \frac{\rho - r_t}{\theta} + c_{t+1} \tag{5.1}$$

(If we make uncertainty explicit, I should technically write this as $c_t = \frac{\rho - r_t}{\theta} + E_t c_{t+1}$.) In the background, I assume there is also a government, whose expenditure is another source of demand. Letting government expenditure, $\log G_t = g_t$, total (log) demand, d_t is then

$$d_t = c_t + g_t = \frac{\rho - r_t}{\theta} + c_{t+1} + g_t \tag{5.2}$$

Idea is, a rise in the interest rate induces the consumer to consume less today (as noted, we assume the substitution effects dominate). But more generally, I could add in other sources of demand such as investment and exports, balance sheet effects etc. These would likely depend on the interest rate too and therefore would reinforce the inverse relationship between demand and the real interest rate. For simplicity, I assume government expenditure is independent of the interest rate.

ASIDE: Relationship to Long-Run Interest Rates

The log Euler equation is

$$c_t = \frac{\rho - r_t}{\theta} + c_{t+1}$$

This implies

$$c_{t+1} = \frac{\rho - r_{t+1}}{\theta} + c_{t+2}$$

Then substituting this into the first Euler equation gives

$$c_t = \frac{\rho - r_t}{\theta} + \frac{\rho - r_{t+1}}{\theta} + c_{t+2}$$

and tidying up

$$c_t = 2\frac{\rho}{\theta} - \frac{1}{\theta}\left(r_t + r_{t+1}\right) + c_{t+2}$$

$$c_t = 2\frac{\rho}{\theta} - \frac{1}{\theta}\left(r_t + r_{t+1}\right) + c_{t+2}$$

More generally, *solving forward n* times gives

$$c_t = N\frac{\rho}{\theta} - \frac{1}{\theta}\left(r_t + \cdots + r_{t+N-1}\right) + c_{t+N}$$

Conveniently, we can use the *expectations hypothesis* to write this in terms of the long-run interest rate. From the expectations hypothesis of the term structure:

$$R^{Nl} = \frac{E_t\left(r_t + \cdots + r_{t+N-1}\right)}{N},$$

where R^{Nl} denotes the interest rate for investing in a N-period long-run bond. Hence,

$$c_t = N\frac{\rho}{\theta} - \frac{N}{\theta}R^{Nl} + c_{t+N}.$$

This makes sense. Consumption today depends on the path of future interest rates and specifically the long-run interest rate. For example, if I expect interest rates to soar in two years time, then that'll tend to reduce my consumption today as I save to exploit this opportunity when it arises. Alternatively, long-run rates will rise today, again inducing a fall in consumption today.

Keep in mind that, long-rates are crucially important, and the bank can, via "expectations management," affect them by changing short rates and expectations of *future* short rates. This gives the power of monetary policy as extra "kick." The above provides a rationale for why *forward guidance* and "expectations management" by the central bank is so important. By committing to keep rates low for a while, they can affect the more important long-run rates. From now on, however, I will simply assume that it is r_t, the short-rate in period t that affects economic activity and thus will continue to use (5.1) as the Euler equation.

5.2.2 Goods Market Equilibrium

Letting y_t denote (log) production, the goods market equilibrium is $y_t = d_t$. So for goods market clearing, production equals demand. The idea here is simple: if demand is 20, production will adjust to meet that demand. In this sense, the goods market "clears." Most importantly, and in contrast to long-run models, we are not even mentioning potential output here. According to this Keynesian short-run analysis, there's no reason whatsoever for $y_t = d_t = y_n$. The *New Keynesian IS curve*—the goods market equilibrium condition—is

$$y_t = \frac{\rho - r_t}{\theta} + c_{t+1} + g_t$$

Most important thing here is the negative relationship between r and y. A lower interest rate leads to a higher level of output/production in equilibrium (in the background, a lower interest rate stimulates consumption, raises aggregate demand, *and then output*.

Now, when output is at its natural rate, the interest rate equals the natural rate. Hence

$$y_n = \frac{\rho - r_n}{\theta} + c_{t+1} + g_t$$

ASIDE: IS Curve in Terms of Output Gaps

Ignore government expenditure for a moment. Then since $y = c + g$ each period, we can write the IS curve as

$$y_t = \frac{\rho - r_t}{\theta} + y_{t+1}$$

In long-run equilibrium, when output is at potential, we have

$$y_n = \frac{\rho - r_n}{\theta} + y_{nt+1},$$

where y_{nt+1} denotes next period's potential output level. Subtracting y_n from y_t then gives

$$y_t - y_n = \frac{1}{\theta}(-r_t + r_n) + y_{t+1} - y_{nt+1}$$

And setting $x_t = y_t - y_n$, we have

$$x_t = \frac{1}{\theta}(-r_t + r_n) + x_{t+1}$$

With this version, we can see clearly how deviations of the interest rate from its natural rate leads to output gaps. This makes it clear that if r_t is set above r_n, the output gap will be negative; i.e., a recession. When the interest rate equals the natural rate, then the output gap is zero. This is why Taylor Rule tries to aim for natural rate *on average*; it is the level consistent with demand equal to *potential*. Recall that in the short-run the economy will not automatically go to natural rate itself—*the* central Keynesian idea.[2]

5.3 The Firm

There are N *monopolistically competitive* firms. I assume N is very large, so the atomistic firm takes aggregate demand Y and the price level P as given. All firms face downwardly sloping demand curves—the firms are not price takers. Yet they take the wage as given. In this model, the interaction of labour supply (by the household) and labour demand (by firms) determines the wage in the labour market; we assume the wage is flexible. The households's labour supply condition implicitly determines its labour supply. The level of production at any given point will determine labour demand; we say there is a "derived demand" for labour. As a result, labour demand rises in booms, while it falls in recessions.

For now, I'm just presenting the technical features of the model. I'll get the role of price stickiness in a moment. As well, I'm dropping the time subscripts, but I should subscript everything with a t below.

The firms faces demand

$$Y_i = \left(\frac{P_i}{P}\right)^{-\eta} \frac{Y}{N}$$

So in a boom Y_i will rise, since aggregate demand, Y, is higher (this could be due to a rise in government expenditure or a fall in interest rates by the FED, say). To be consistent with the analysis before, the Y I put here should really be the d_t in Eq. 5.2 above. But forget about the distinction between logs of variables and their actual levels for now—this distinction has no bearing on the central ideas.

Now, by choosing P_i, the firm implicitly chooses Y_i too. Note that P is the price *level* in the economy, which you can think of as simply the average price set by all firms. Formally, it would correspond to a price index such as the CPI.

What determines demand for the firms product is its *relative price*, not its absolute price P_i. This is what the firm will keep in mind and what we will ultimately try to solve for. Moreover, by setting a relative price, it ensures it's maximizing *real* profits—what the firm actually cares about. For example, setting a price of 100 in an environment were all other prices are around 1000000 will obviously not maximize *real profits*. Notice that, even if a firm sets a relatively high price, there is still positive demand for their product. We often

[2]For instance, in the long-run classical model, if consumption fell, then the rise in savings would cause the natural interest rate to fall, which would cause investment to rise. And that induced investment would rise to clear the goods market, so we'd end up back at potential again with lower consumption, but higher investment. By contrast, Keynes argued that if consumption fell, *output* (and to a much less extent, the interest rate) would adjust and so, instead, the economy would enter a *recession.*

explain this by saying the consumer has a "love of variety" and would like to purchase a little of everything. Think about it; when you buy yoghurts, you might purchase one of each flavor even if some are more expensive than others. At a micro level, this "love of variety" can be rationalized by diminishing marginal utility to individual *goods*; therefore, you'd rather have a strawberry and apple yoghurt rather than two apple ones. See? The elasticity of demand, η, depends on the substitutability between goods. As we'll see in a moment, this will determine the markup.

The firm's production function is

$$f(L) = L,$$

where L is number of workers hired by firm. In this version, there is no capital or investment. Note that the marginal product of labour is $MPL = 1$. (To make things more realistic, we could also have $f(L) = AL$, where A denotes productivity.) The production function will determine the marginal cost in this model. Because $A = 1$ and labour is the only factor of production in our basic model, real wages will ultimately determine marginal cost. It follows that if a firm wants to produce Y_i units, it needs to hire Y_i workers. As I said, there is a *derived demand* for labour: it derives fundamentally from the level of aggregate demand.

Now, if a firm charges P_i, the demand for its goods is $\left(\frac{P_i}{P}\right)^{-\eta}\frac{Y}{N}$. The firm's revenue is then $P_i\left(\frac{P_i}{P}\right)^{-\eta}\frac{Y}{N}$. Its labour demand will be then Y_i, and its costs are WY_i. Overall, by setting a price of P_i, the firm's profit is then

$$\Pi_i = P_iY_i - WY_i = Y_i(P_i - W)$$

Substituting in the expression for demand gives

$$\Pi_i = \left(\frac{P_i}{P}\right)^{-\eta}\frac{Y}{N}(P_i - W)$$

The only choice variable is P_i. As noted above, the wage is exogenous to firm, and will depend on developments in the national labour market (in particular, firms' labour demands interacting with household's labour supply).

Maximizing with respect to P_i is[3]

$$P_i = \frac{\eta}{\eta - 1}W$$

[3]Regarding second order conditions, $\Pi_i'' < 0$ (i.e., the second order derivative with respect to P_i), so this is indeed a maximum.

In this world, marginal cost equals the wage, so more generally, $P_i = \frac{\eta}{\eta-1}MC$. (It's typical to write this result in terms of marginal cost, so, later on, if I replace the formula above with $P_i = \frac{\eta}{\eta-1}MC$, don't be alarmed.) This is what the firm would charge in a *flexible price equilibrium*. The constant, $\frac{\eta}{\eta-1}$ is the firm's target markup. <u>No matter what</u>, the firm will always seek this markup. At any point in time, this is the ideal price, *no matter what the current markup or level of demand is*. This is a key point to consider, especially when we deviate from equilibrium. Note too that the markup is pinned down by a fundamental, structural feature of the economy; i.e., the elasticity of demand.

There are a few points to note here. Most significantly, the firm charges a price that exceeds the marginal cost of production. This is a consequence of the firm's monopoly power. You see, to maximize profits, it considers the "menu of options" given by the demand curve; for example, it could charge a high price and have low demand/output; or a low price and high demand/output. Ultimately—as the maths tells us—what maximizes its profits is setting a price of $P_i = \frac{\eta}{\eta-1}W$ and producing $Y_i = \left(\frac{\frac{\eta}{\eta-1}W}{P}\right)^{-\eta}\frac{Y}{N}$. (To get the latter, just substitute the firm's optimal price into its demand curve.) To give a concrete example, suppose the maths tells us that a firms optimal price is 10, and its optimal quantity is 30. Suppose its constant marginal cost is 8. Of course, the maths tells us that producing more will *reduce profits*. Why? Well, if the firm wanted to produce 31, say, it'd have to reduce its price to 9, say. So, it gains an extra 9, but *loses* 1 on all existing units. In this case, therefore, its marginal revenue from producing an extra unit would be $9 - 30(1) = -21$. Hardly an attractive option, is it?

Because $P > MC$ in monopolistic competition, output is below the socially optimal level. An implication of this is that, *all else constant*, people would be happy to pay P' where $MC < P' < P$, say, and both the firm and these people would be made better off. Yet this transaction never occurs. Technically, the monopolistically competitive equilibrium is not Pareto optimal, and the First Welfare Theorem does not hold. This is an important consequence of monopoly power. By contrast, in the ideal market structure of perfect competition, the firm always produces where $P = MC$.

Moving on, the firm is mainly concerned with is relative price, since that is what determines its demand. To get the firm's desired relative price, just divide the formula above by the price level P:

$$\frac{P_i}{P} = \frac{\eta}{\eta-1} = \frac{\eta}{\eta-1}\frac{MC}{P}$$

This is firms *ideal* relative price if it were free to adjust. So if the real wage, $\frac{W}{P}$, rose, the firms desire price, P_i, would also rise.

5.4 Aside: Equilibrium

What does an equilibrium look like in this economy? From above we have

$$\frac{P_i}{P} = \frac{\eta}{\eta - 1}\frac{W}{P}$$

Now, in equilibrium, we assume all firms are the same, and face the same marginal costs. In this *symmetric equilibrium*, all firms will charge the same price, so $P_i = P$. So, in equilibrium, each firm has a relative price of one. Moreover, each firm's price equals the price level. If all students in class set a price of 10, then the price level in the class will obviously be 10. Thus in a symmetric equilibrium,

$$\frac{P_i}{P} = 1 = \frac{\eta}{\eta - 1}\frac{W}{P}$$

This implies the equilibrium real wage is

$$\frac{W}{P} = \frac{\eta - 1}{\eta}$$

Note in particular that this is less than 1, and 1 is the marginal product of labour. The fact that the firm pays the worker less than what he produces—i.e., $\frac{\eta-1}{\eta} < 1 = MPL$—is of course the source of his profits (and is another way of saying price exceeds marginal cost). The firm makes a profit of $1 - \frac{\eta-1}{\eta}$ per unit. For instance, if $\eta = 3$, then the real wage will be $\frac{2}{3}$; so the firm will earn $\frac{1}{3}$ on each unit produced.

Now, equilibrium production by any firm is then $\left(\frac{\frac{\eta}{\eta-1}W}{P}\right)^{-\eta}\frac{Y}{N}$ So with an equilibrium real wage of $\frac{W}{P} = \frac{\eta-1}{\eta}$, equilibrium production by any firm is $\frac{Y}{N}$. Hence, by symmetry, total equilibrium production in the economy will by Y. This will be our Y_n. Note that this is also equal to equilibrium aggregate demand.

So what is equilibrium output, anyway? Well, production was entirely determined by labour; this was the only factor of production. There's simply no other way to get more output, given the production function. So, fundamentally, we must ask how much is the household willing to supply at the equilibrium real wage of $\frac{W}{P} = \frac{\eta-1}{\eta}$. To determine this, we go to the labour supply optimality condition; i.e., the first order condition for labour. And then just find out what labour supply is at this equilibrium wage, $\frac{\eta-1}{\eta}$. We can get this from the labour supply curve. I won't go into details here, but just think of picking the point on the labour supply curve where the real wage is $\frac{\eta-1}{\eta}$.[4] This will pin down

[4] We know therefore that, in equilibrium, $\frac{\eta-1}{\eta}\frac{1}{C^\theta} = L^\sigma$. But labour supply must equal output, so $\frac{\eta-1}{\eta}\frac{1}{C^\theta} = Y^\sigma$.

equilibrium labour supply and hence equilibrium output/production. This level will be the level of *potential or natural* output. This is determined by fundamentals and the interaction of household/firm maximization. Because of monopolistic competition and the fact $P > MC$, though, output will be below its socially efficient level.

That's it. We now have equilibrium real wage, equilibrium output/production/demand and equilibrium labour supply. All actions are consistent. Labour demand equals labour supply. Aggregate demand equals the level of production. We haven't look at the breakdown of output between consumption and government expenditure, but we know that the goods market will clear: $Y = C + G$. To get the natural rate of interest, we go back to the Euler equation $u'(C_t) = \beta(1 + r_t)u'(C_{t+1})$. In equilibrium, we know that $C = Y - G$, where Y is the natural rate of output and G is some constant. Assume too that the natural rate of output is constant, which must be true here, since there's no obvious way for it to rise over time.[5] Substitution $C = Y - G$ into the Euler equation and noting that "nothing changes" in an equilibrium gives

$$u'(Y - G) = (1 + r_n)\beta u'(Y - G) \Rightarrow r_n = \frac{1}{\beta} - 1.$$

If we introduced money, the usual money market equilibrium condition must be satisfied in equilibrium[6]

$$\frac{M}{P} = L(r_n, Y_n)$$

Here Y_n and r_n have already been determined. In this model, if prices are flexible, a rise in M just rises P, and *money is still neutral.* Notice that P must adjust since r_n and Y_n have already been nailed down by fundamentals.

5.5 Increase in Demand at Potential

We now turn to see how output can be demand determined in the short run. Central to this story is price rigidity. For now, I simply <u>assume</u> prices are fixed (by price controls, say). With prices fixed, we will see how "money matters." Note, however, that the increase in demand, could be attributable to anything—such as a rise in government expenditure, less

Now, if we assume that consumption C is some fraction γ of output, Y, in equilibrium (so government expenditure makes up a fraction $1 - \gamma$), then solving $\frac{\eta-1}{\eta}\frac{1}{(\gamma Y)^{\theta}} = Y^{\sigma}$, will give us equilibrium output/labour supply.

[5] If the production function was $Y = AL$ and technology, A, was growing, then the natural rate of output would rise over time.

[6] To generate a demand for money like this, we'd have to put money in the utility function, but rest assured, we can easily do that and derive a money demand from microfoundations.

precautionary savings, and so on. For now, I'll focus on the case where the money supply increases, which causes a fall in the real interest rate, which then stimulates an increase in aggregate demand to $Y > Y_n$. To get the intuition, it might be convenient to simply thing of a "helicopter drop" of money, which raises demand to Y. Anyway, demand is now

$$Y_i = \left(\frac{P_i}{P}\right)^{-\eta} \frac{Y}{N}$$

Each firm now faces increase in demand. To produce more, labour demand increases nationally. From the production function, we know that this is the only way the firm can produce more. Yet to induce greater labour supply, real wages must now rise. Given C_t the labour/leisure optimality condition implicitly defines a labour supply function:

$$L_t = \left(\frac{W_t}{P_t}\frac{1}{C_t^{\theta}}\right)^{\frac{1}{\sigma}}$$

By the permanent income hypothesis (i.e., the Euler equation), any temporary increase in the wage should be smoothed over the entire lifetime. So the effect on consumption today of a temporary increase in the wage will be small. After all, households know this is merely a *transitory* increase as a result of the business cycle. Therefore, it's fine to treat the above as an increasing relationship between the real wage and labour supply (assuming the real wage increase is transitory and therefore doesn't affect C too much.) Point is, to increase L, $\frac{W_t}{P_t}$ must rise. By how much? This depends on σ: how painful are those extra hours of labour?

Clearly this is not good for the firm. Ideally, now, they'd like to raise price, which you'll recall is increasing in marginal cost—and *that* has just shot up. Specifically firm i aims for:

$$\frac{P_i}{P} = \frac{\eta}{\eta - 1}\frac{W}{P}$$

But because P_i is fixed, and $\frac{W}{P}$ has <u>risen</u>, the markup now deviates from desired one, $\frac{\eta}{\eta-1}$.[7] In particular, markup is now less than the desired markup $\frac{\eta}{\eta-1}$. This is not good for the firm! Remember, they always desire the target markup $\frac{\eta}{\eta-1}$.

Because the price is fixed and the real wage has risen, the markup falls. Given firm's profit maximization objective, this suboptimal markup is unsustainable. If they desired this lousy markup, they'd have charged that at equilibrium. But they didn't. Yet, given the price control, what is the best thing that a firm can do? Well given that $P > MC$, the best thing they can do is <u>meet the demand</u>; they are, after all, earning *some* markup selling more

[7] But more generally, for *any* arbitrary price P_j and marginal cost $\frac{W}{P}$, $\frac{P_j}{P} = \text{MARKUP}\frac{W}{P}$.

units. *This is the best thing they can do in the circumstances.* But for this to happen, we must assume the real wage doesn't jump up *too much*. The firm will only be happy to increase output as long as $P > MC$. If the marginal cost jumps up to $MC > P$, then they will not meet the increase in demand.

Eventually (in the proverbial "long-run"), when the firms do get an opportunity to change prices, *they will.* They do want their target markups back. At this point, all firms start raising prices, causing the *price level*, P, to rise. In turn, this causes the purchasing power of the money supply to fall, which reduces demand. (In practice, the incipient rise in prices would induce the central bank to start raising interest rates and withdrawing the money from the economy.)

Money Market Equilibrium

In the New Keynesian model, the money market is not explicit. So where is the money, then? Well, implicit in the Taylor rule is the idea that the central bank is changing the money supply to hit its target; this way, it is countering money demand shifts behind the scenes.

Nonetheless, it is convenient to also view dynamics through the lens of the money supply/demand analysis. From the money market equilibrium condition, $\frac{M}{P} = L(r, y)$, we see that, with prices initially fixed, the rise in the money supply will case r to fall and y to rise. Conversely, when prices start to rise, this causes the real interest rate to rise and output to fall. (Somewhat more intuitively, you could say that if some firms raise their prices, then that'll leave their customers will less money to spend on *other* firms, meaning the level of real aggregate demand will fall.) As a result, labour demand and real wages fall. This process will continue until we get back to potential, and ultimately the price level will rise in proportion to the rise in the money supply. Nothing *real* will change.

Comments

Rather that saying "money is nonneutral" in the short run, we know have a story, a mechanism, through which this happens. It shows how the FED can temporarily affect output. Moreover, we also have the "money is neutral in the long run" story too. We've seen how the economy adjusts back to the potential equilibrium; this reversion to the natural (or potential) rate of output called the *natural rate hypothesis*. In particular, we've seen how starting at

$$\frac{M}{P} = L(r_n, Y_n)$$

an increase in M with P fixed causes Y to rise above Y_n.

Exactly why this happens should be clear. Central to the story are sticky prices; formally they are the *propagation mechanism* of the model. This is consistent with the idea that the FED can affect affect output temporarily, but not forever. As long as prices are sticky, the FED can stimulate economy. Thus booms and busts can be quite persistent, especially if prices adjust sluggishly.

Note that when demand increases, production increases too to meet that level. Most importantly, we have rationalized the idea of *demand-determined* output. We've seen how changes in demand induce proportional changes in output. And just to be clear, the increase in demand could have been generated by changes in government expenditure or changes in consumption (or, in a more general setting, investment, exports, and so on.) Of course, to get this, we assume marginal cost doesn't rise above the preset price.[8]

Imagine now if the money supply increases and prices are flexible. In a *flexible price* economy, we assume the above adjustment process happens in an instant. The increase in labour demand instantly pushes up real wages, and firms instantly increase their prices and demand falls back to potential.[9]

Recessions

Content yourself that the opposite situation occurs when the money supply falls (or, for that matter, if consumption demand falls). If say consumption expenditure suddenly fell, then demand would fall and firms, keeping prices fixed, would simply cut back production. In turn, labour demand would fall, along with real wages. In this case, then, the firms keep their prices *too high.* As above, eventually, they'll change their prices to a lower level, which will increase aggregate demand again. Of course, this kind of story provides a rationale for the FED to lower rates in recession. Point is, the lower interest rates would raise consumption demand and therefore raise aggregate demand (and hence production) again. The Keynesian recommendation was for the government to increase expenditure to make up the short-fall in demand; that is, the government, via fiscal policy, should increase expenditure in recessions.

[8] Note that we must use imperfect competition to model this. If $P = MC$ and marginal cost rises, then the firm will *not* meet the demand; otherwise it would suffer losses. For this reason, to model money nonneutrality, we need some form of imperfect competition.

[9] More realistically, firms would pre-empt the rise in the price level and they would adjust prices immediately; labour demand wouldn't rise at all.

A Note on the Labour Market

The labour market is a little strange in the model, but fortunately that has no bearing on anything important. The way I've modelled it, we just have a single household or agent supplying labour hours. Implicitly, I've assumed the household spreads those working hours and bodies over all firms. Most notably, the labour market clears. Because of this, there is no unemployment in the model; the household simply increases and decreases labour market activity as the firm "seduces" it into doing so by changing wages. That is, it's just the same people reducing and lowering hours worked all the time. There are no new people entering the workforce. You might suspect that labour markets are rather different in reality. However, none of this really matters for what the model seeks to address.

5.5.1 Real and Nominal Rigidity

A crucial question for the model is, why are prices so sticky anyway? If prices are flexible, the the demand-determined output prediction falls apart. There are two ways to rationalize sticky prices. The first is *nominal rigidity*. This refers literally to some "menu costs" of changing prices. Realistically, however, these would have to be fairly large. If marginal costs rise sharply in booms, then the incentive for firms to raise prices is surely large and might well overwhelm any "menu cost." Therefore, although we need some nominal rigidity, we need what's called *real rigidity* too.

What is real rigidity? Real rigidity measures the degree to which marginal costs change due to fluctuations in output. If there was real rigidity in an economy, there would be little change in firms' real marginal costs—the source of price changes—over the business cycle. To generate real rigidity, we need a way to rationalize marginal cost stability. In turn, this would generate price rigidity even in the face of (realistically) small "menu costs." For this reason, a greater degree of real rigidity leads to greater money nonneutrality.

The essential idea behind real rigidity is this: is there something that prevents marginal costs from rising a lot in booms (or conversely, prevents marginal costs from falling a lot in recessions)? Throughout, and without loss of generality, I will stick to the boom case. Before going on, recall the formula for the firm's optimal price at any point:

$$\frac{P_i}{P} = \frac{\eta}{\eta - 1}\frac{MC}{P},$$

where $\frac{MC}{P}$ denotes the firms real marginal cost. Note further, that when we have the production function, $Y = AL$, the marginal cost is $\frac{W}{MPL} = \frac{W}{A}$. While by no means exhaustive, here are some sources of real rigidity:

1. *Increasing returns to scale* or some kind of "learning by doing." With such features, productivity would rise in a boom. So if we have the real wage rising, and productivity rising simultaneously, then real marginal costs won't rise as much (and could in fact fall!) Certainly, this reduces the firms' incentives to change their prices. (By contrast, in a recession, this will cause productivity to *fall*. In turn, this *reduces* the firm's incentives to *lower* its price.)

2. *Implicit Contracts*: Say firms make a deal with workers to keep their real wages stable over time. Then, their wages won't rise in booms and won't fall in recessions. If workers have little access to capital markets, then this will act as an insurance for them, and help them smooth their consumption over time. On the other hand, this would make firms profits more variable. This kind of real rigidity could be rationalized by assuming firms are less risk averse than workers.

3. *Balance sheet effects.* If property and equity prices rise in booms, then this will mean firms can offer banks more collateral for loans. Because the value of their collateral has gone up, this means a) banks will be more willing to lend and b) banks will charge the firms a lower risk premium in lending to them (formally, the firm's "external finance premium" falls in booms). Empirically, risk premia on corporate bonds fall in booms. Overall, this lowers firm's real marginal cost, attenuating their desire to raise prices.

4. *Cyclical Variation in Markups*: If the firm's elasticity of demand rises in booms due to, say, to greater competition or firm entry, then their desired markup will fall. The interaction of a rising real wage and a lower desired markup would, on net, lead to less upward pressure on prices.

5. *Shifts in the labour supply curve.* If labour supply increases in a boom for any given level of the wage, that that will mitigate any increase in the equilibrium real wage. Immigration, for example, would have this effect.

Note that if there were, say, trade unions who increased wage demand in a boom (due to more bargaining power as a result of "tighter" labour markets), then there would be more upward pressure on real wages. Of course, this would *lower* the degree of real rigidity, making it harder to rationalize sticky prices. With such a feature, prices would surely be more flexible, attenuating the power of monetary policy.

Note finally why we need some nominal rigidity too. If we didn't have nominal rigidity, prices would still change—unless the degree of real rigidity miraculously caused marginal costs to stay *constant*. Think of it like this: real rigidity will cause marginal costs not to rise as much as they *"should,"* and then the nominal rigidity or menu cost makes it optimal for the firm to leave prices as they are.

5.6 Price Setting in a Sticky Price Environment

Now, we want to formalize price setting. In the basic New Keynesian model, I've assumed all prices are fixed in the short-run. Yet, in reality, some firms are always adjusting, while others are not. This does not affect the fact that money is nonneutral, since the *price level* is still sticky. But it *does* mean that the price level will adjust somewhat to today's output gap. The New Keynesian Phillips curve gives a description of how the price level and hence inflation changes over time. To give the intuition, I'll give a baby example first.

From above, the optimal pricing rule for firm i is

$$\frac{P_i}{P} = \frac{\eta}{\eta - 1}\frac{W}{P}$$

For now, just assume the marginal cost comes completely from real wages, as it did above in the model.

Now, take logs to get

$$\log P_i - \log P = \log \frac{\eta}{\eta - 1} + \log \frac{W}{P}$$

As is standard, write lower case letters for logs; i.e., $\log X = x$ etc:

$$p_{it} = p_t + \log \frac{\eta}{\eta - 1} + \log \frac{W}{P}$$

To give a baby example, suppose firm i sets prices *today* for *two* periods, so they are "locking" themselves in now. (In the background, imagine there is some "menu cost" to changing prices, making it optimal for the firm to keep prices fixed for 2 years.) Now, the optimal price in period 1 for firm i is:

$$p_{i1} = p_1 + \log \frac{\eta}{1 - \eta} + \log \left(\frac{W}{P}\right)_1$$

Firm i's expected optimal price next year is

$$E_1 p_{i2} = E_1 p_2 + \log \frac{\eta}{\eta - 1} + E_1 \log \left(\frac{W}{P}\right)_2$$

where E_1 denotes expectation as of time 1. An obvious solution to the firm's problem is to set the price today equal to an average. That is, the forward-looking firm—who has rational expectations—sets a price today of

$$p^* = \frac{p_{i1} + E_1 p_{i2}}{2}$$

This way, we come as close as possible to maximizing profits each periods. Note, however, prices are now suboptimal each period.
Substituting gives the price the firm will set today:

$$p^* = \frac{1}{2}\left(p_1 + \log\frac{\eta}{1-\eta} + \log\left(\frac{W}{P}\right)_1\right.$$

$$\left. + E_1 p_2 + \log\frac{\eta}{\eta-1} + E_1 \log\left(\frac{W}{P}\right)_2 \right)$$

More generally, we could extend this to arbitrarily many periods. (If the firm cared less about future profits, we'd have a weighted average, with more weight on the present; but forget about this for now.) There is a lot of insight here already. First, the price set today depends on expected future real marginal costs. For example, if the firm expects high real wages next period, that will raise *today's* price (and hence price *level*). Second, the price set today will depend on the expected *future* price level, $E_1 p_2$; the firm cares about its *real profits* and *relative* price. Thus, if for some reason, the firm expects higher prices on average in the economy next period, that will raise the firm's optimal price *today.* For example, if the firm predicts a large depression next period, the firm might predict other firms—who might be free to adjust prices—will lower their prices next period, thereby causing a fall in the *price level*. Likewise, if the central bank commits to a high price target next period, that would induce the firm to set a higher price this period. Hence expectations of future monetary policy affects price-setting behaviour *today*.

5.7 Price Dynamics and the Business Cycle

From now on, it's useful to think of a more general production function where MC could incorporate a range of inputs such as oil or the cost of capital. To start with, consider the firm's optimal relative price

$$\frac{P_i}{P} = \frac{\eta}{\eta-1}\frac{MC}{P}$$

Taking logs and imposing the equilibrium condition, $P_i = P$:

$$0 = \log \frac{\eta}{\eta - 1} + \log \frac{MC}{P}$$

Using log rules

$$\log \frac{MC}{P} = \log \frac{\eta - 1}{\eta}$$

Denoting the equilibrium log real marginal cost by $\widetilde{\log \frac{MC}{P}}$, we have $\log \frac{\eta-1}{\eta} = \widetilde{\log \frac{MC}{P}}$
More generally, we have

$$\frac{P_i}{P} = \frac{\eta}{\eta - 1} \frac{MC}{P}$$

$$\log P_i - \log P = \log \frac{\eta}{\eta - 1} + \log \frac{MC}{P}$$

Using log rules again:

$$p_{it} = p_t - \log \frac{\eta - 1}{\eta} + \log \frac{MC}{P}$$

But from above we know that $\log \frac{\eta-1}{\eta} = \widetilde{\log \frac{MC}{P}}$. Substituting then gives

$$p_{it} = p_t + \log \tfrac{MC}{P} - \widetilde{\log \tfrac{MC}{P}}$$

$$p_{it} = p_t + \log \tfrac{MC}{P} - \widetilde{\log \tfrac{MC}{P}}$$

This makes sense; your target price deviates from equilibrium if real marginal cost deviates from equilibrium. In practice, and as should be clear from the basic model above, $\log \frac{MC}{P} - \widetilde{\log \frac{MC}{P}}$ is proportional to output gap, $y_t - y_n$. That is, for some $\alpha > 0$:

$$\log \tfrac{MC}{P} - \widetilde{\log \tfrac{MC}{P}} = \alpha(y_t - y_n)$$

In other words, when output goes above potential, real marginal cost rises above its equilibrium level (i.e., the level at potential). It follows, therefore, that the firms optimal (log) price at any time is:

$$p_{it} = p_t + \alpha(y_t - y_n)$$

If the firm were free to change prices, *this* is what it would change to. Several points are worth noting here. The parameter, α mediates the degree to which marginal cost responds to output. Because changes in marginal cost are the underlying sources of price level changes (and hence inflation) in the model, deviations in output from potential play a central role in pricing pressure. In this sense, α partly mediates the degree of *real rigidity* in the economy: how responsive is price to deviations of output from potential? For example, if α is low, then there is a lot of real rigidity; prices don't change much as output varies over the business cycle. Of course, a low α could also reflect a high degree of nominal rigidity too. Keep in mind, that the extent to which wages—and hence marginal costs—rises depends on such factors as the elasticity of labor supply. Point is, the parameter α depends on features of the economy.

5.8 The New Keynesian Phillips Curve

We assume now that *some* firms change prices each period. Opportunities to change prices are *time-dependent*.[10] In particular, each firm faces some fixed probability of changing price each period. For example, if that probability is .1, then there will, on average, be little opportunity to change prices. On average, the firm's price in this case will be fixed for ten periods. Crucially, the forward looking firms—with rational expectations—take account of this when the have an opportunity to change. For example, if this is a .1 world, they'll set prices keeping in mind developments far into the future. Namely, when they change their prices they know they're "locking themselves in" for a long time. Therefore, when they're setting prices today, they'll put a good deal of weight on the future optimal prices, since they know they mightn't get another chance to reset for a while. So they'd better get in right and pay a lot of attention to the future when setting prices *today.* This formulation of modelling price setting is attributable to Guillermo Calvo.

To summarize, when firms get an opportunity to change prices, they'll consider a) optimal prices in future b) probability of changing again. If there's a greater chance of an opportunity to change prices in the future, they'll place less weight on future prices when they setting prices today. Rather than placing much weight on future prices *today*, it would be better to wait until the future, when they're likely to get another chance to change. But if opportunities to change are rare, the firm *will* place more weight on the future when they're setting prices today.

As shown above, the firm's optimal (log) price at any time t is

[10] With *state-dependent* pricing, firms change prices depending on the state of the economy.

$$p_{it}^* = p_t + \alpha(y_t - y_n)$$

Letting $x_t = y_t - y_n$, we have

$$p_{it}^* = p_t + \alpha x_t$$

Having received a tap from the Calvo fairy, a fraction δ of firms have an opportunity to change their price in any given period. This is also the probability of getting a chance to change each period; it measures the degree of *nominal* rigidity (recall that *real* rigidity measures the *extent* to which firms want to change their prices.) So a fraction $1 - \delta$ leave things be; they don't get a chance to change. Note that the price level in a given period is the average of all prices in the economy; here it is a weighted average of those who change and those who keep their stale prices from last period. So the (log) price level is

$$p_t = \delta p_t^* + (1 - \delta) p_{t-1}.$$

Note now that changes in logs are equal to time derivatives.[11] Then, taking $\log P_{t-1} = p_{t-1}$ from each side of the above gives

$$\pi_t = \delta \pi_t^*. \tag{5.3}$$

We'll come back to this in a moment. Recall the optimal price formula from above: any firm changing prices in period t will charge

$$p_t^* = p_t + \alpha x_t$$

Take p_{t-1} from each side to get

$$\pi_t^* = \pi_t + \alpha x_t$$

By choosing prices, firms are implicitly choosing inflation rates. That's all this means. So optimal prices implicitly give optimal inflation rates. Rather than thinking of firms setting prices, we think of firms setting inflation rates.

Now consider a firm changing its price today and locking itself in. What does it do? Analogous to the two-period example above, the firms sets the current prices as a weighted average of all the future optimal prices. The weights depend on two factors. First, how

[11] By the chain rule $\log P_t - \log P_{t-1} \approx \frac{d \log P_t}{dt} = \frac{d \log P_t}{dP_t} \frac{dP_t}{dt} = \frac{\frac{dP_t}{dt}}{P_t}$.

much does the firm care about future profits? This will depend on the firm's discount factor, $\phi < 1$. Second, what are the chances the firm will get to reset prices in the future? Obviously, if the firm has lots of chances to reset (i.e., δ is high), then it'll place less weight on expected optimal future prices (since it correctly assumes it'll surely get an opportunity to change soon). That's why the $(1-\delta)$s appears in front of them. Thus, for example, if opportunities to change come frequently (i.e., δ is high), the firm will downweight future prices highly by placing a $1-\delta$ in front of future optimal prices. Accounting for all these factors, when the firm gets an opportunity to change, it sets a price of

$$\pi_t^* = \pi_t + \alpha x_t + (1-\delta)\phi(E_t(\pi_{t+1} + \alpha x_{t+1}) +$$

$$+(1-\delta)^2\phi^2(E_t(\pi_{t+2} + \alpha x_{t+2}) + (1-\delta)^3\phi^3(E_t(\pi_{t+3} + \alpha x_{t+3}) + ...,$$

where this summation is infinitely long. Note that if $\delta = 1$, then firms will only take account of today's output gap when setting prices today. This makes sense; it'll have another chance to change next period and will be able to respond best to next period's developments. The ϕs just capture the fact that the firm cares less about future profits as ϕ falls; hence it'll pay less attention to "getting things right" in the future when its changing prices today. Instead, it cares more about "getting it right" *today* and thereby maximizing today's profits. Now from (5.3) above we know that:

$$\pi_t = \delta\pi_t^*$$

Substituting our expression for the optimal price into the above gives

$$\pi_t = \delta\left(\pi_t + \alpha x_t + (1-\delta)\phi(E_t(\pi_{t+1} + \alpha x_{t+1}) + (1-\delta)^2\phi^2 E_t(\pi_{t+2} + \alpha x_{t+2}) + ...\right) \quad (5.4)$$

And from this, get $E_t\pi_{t+1}$ and hence $(1-\delta)\phi E_t\pi_{t+1}$:

$$(1-\delta)\phi E_t\pi_{t+1} = (1-\delta)\phi\delta\Big(E_t(\pi_{t+1} + \alpha x_{t+1}) + (1-\delta)\phi E_{t+1}(\pi_{t+2} + \alpha x_{t+2}) + ...\Big) \quad (5.5)$$

Now, subtract the above from (5.4) to get

$$\pi_t - (1-\delta)\phi E_t\pi_{t+1} = \delta\pi_t + \delta\alpha x_t$$

Tidying this up yields

$$(1-\delta)\pi_t = (1-\delta)\phi E_t \pi_{t+1} + \delta\alpha x_t$$

$$\Rightarrow \pi_t = \frac{\alpha\delta}{1-\delta} x_t + \phi E_t \pi_{t+1}$$

Then writing the output gap more formally:

$$\pi_t = \frac{\alpha\delta}{1-\delta}(y_t - y_n) + \phi E_t \pi_{t+1}$$

This is the *New Keynesian Phillips Curve.* Really, this is based on reasoning in two-period example above. First, today's inflation depends on today's output gap, $y_t - y_n$. Recall how the output gap was a proxy for the deviation of *real* marginal cost from its equilibrium value. So a high output gap is coincident with high marginal cost. That is, a large output gap today will lead to upward price pressure since marginal costs are rising. Notice how α appears; as we know, a high α makes marginal cost more sensitive to output gaps and will lead to greater upward pressure on prices for any *given* output gap. Second, today's inflation will depend on today's expectation of future inflation.[12] As noted above, those setting prices today consider what happens to the price level in future periods. For example, if they expect a high price level next period, to maintain real profits—as they "lock" themselves in—they will set a higher price today. Equivalently, if they expect high inflation next period, they will set high prices *this* period, leading to high inflation *today*.

Note that as δ rises, meaning there are more firms changing price each period, the coefficient on the output gap rises. This makes sense: as the proportion of firms changing today increases, price increases will respond more to today's output gap. If firms are changing prices more often, then there will be more inflation pressure for any given output gap. Realistically, too, δ surely depends on the rate of inflation (by the Lucas Critique).

Finally, don't confuse *disinflation* and *deflation*. Deflation is when prices actually *fall*; a sufficiently large output gap will lead to deflation. Disinflation, by contrast, is a reduction in the rate of inflation, say from 15% to 10%. So, under *disinflation*, we can typically have prices still *rising*. Think of it like this: deflation is walking backwards, disinflation is slowing down, and hyperinflation is sprinting.

[12] Note that inflation next period, π_{t+1}, depends on marginal costs *next* period and expected inflation in the *following* period. The latter in turn depends on marginal costs in *that* period. Using this reasoning, inflation today depends on all future expected marginal costs.

5.9 Three Equation Model

Writing $r_t = i_t - E_t\pi_{t+1}$, we have derived the New Keynesian three-equation model:

$$y_t = \frac{\rho}{\theta} - \frac{1}{\theta}(i_t - E_t\pi_{t+1}) + \mathbb{E}c_{t+1} + g_t + v_t$$

$$i_t = r_n + \pi_t + \gamma(y_t - y_n) + \beta(\pi_t - \overline{\pi}) + \varepsilon_t.$$

$$\pi_t = \phi E_t\pi_{t+1} + \zeta(y_t - y_n) + u_t$$

The terms v_t, ε_t, and u_t capture movements in y_t, i_t and π_t that are unrelated to what is already included in the equations. γ and β represent the weights the bank places on output and inflation, respectively. If the bank cares little about output, then $\gamma \approx 0$.
To account for interest rate smoothing the Taylor rule is sometimes written

$$i_t = \rho i_{t-1} + (1-\rho)\left(r_n + \pi_t + \gamma(y_t - y_n) + \beta(\pi - \overline{\pi})\right).$$

That is, the optimal interest rate is a weighted average of last period's rate and the optimal one dictated by the standard Taylor rule *this* period.

The early eighties in U.S. provides a useful example of these three equations in action. In the early eighties, Paul Volcker, the then FED Chair raised nominal interest rates to almost 20% to restrain inflation, leading to an enormous recession around 1981. Unemployment subsequently rose to about 11%, and output fell well below the natural rate. However, inflation fell from about 10% to 4% by around 1986 and output and unemployment had by then returned to their natural rates. The disinflationary policy worked. Note, however, that prices were rising all the time, but there was a *moderation* in the rate of wage and price increases.

Chapter 6

Real Business Cycle Theory

6.1 Introduction

The emphasis here is on technology/TFP shocks, and the associated supply-side responses. As the term suggests, all the shocks are real, not nominal. As we know from the Solow model, total factor productivity, A, is central to maintaining sustained economic growth. But if A is so important in the long-run, then surely, isn't it likely A will be important in the short run too? This is the motivation for this theory. Recall from the Solow model how we assumed A grew at some constant rate, say 2%, in the long run. But why should the growth of A be so smooth? Imagine A starts growing above trend, say at 4%, for a couple of years. How would this affect the economy? This is the question RBC theory addresses. Remarkably, temporary but somewhat persistent fluctuations in technology, A, *lead to exactly the kind of business cycles we see in the data*. The real business cycle model can replicate business cycle fluctuations without *any* reference to demand, Keynes, sticky prices, or the money supply. From a modelling perspective alone, this is impressive.

6.1.1 Introduction

As in the New Keynesian model, I will restrict myself to the case of a boom, but, by symmetry, the opposite occurs in a recession. For attaining intuition, it is useful to think of A being constant at potential and then rising for a few periods.

In the model, the constant returns to scale production function is Cobb-Douglas:

$$Y = A_t K_t^{\alpha} L_t^{1-\alpha}.$$

The main player in the model is technology, A. Yet A should be interpreted broadly. It is anything that changes the level of production, Y, for a given K and L. For instance, it could incorporate inventions, oil shocks, weather, regulation, and so on. Another relevant example is a shock to the financial system: this would reduce the efficiency of the economy in allocating resources to their most efficient uses and, in aggregate, this would act as a fall in A. Formally, A is a random variable that follows an autoregressive stochastic process: A is related to its previous value, but there is some random element added on. This causes shocks to A to be *persistent*; i.e., a high A today will lead on average to a high A tomorrow. For example, ignore trends, and suppose the steady state value of A might be 1. Then this period A might rise to 1.5; then next period it will be 1.3, then 1.2, and so on; after a few periods A will revert to its steady state value of 1 again.

Ignoring government and net exports, in the model we have

$$Y_t = A_t K_t^{\alpha} L_t^{1-\alpha} = C_t + I_t.$$

Here, the sources of demand are C_t and I_t. As A_t varies, so does potential Y_t. In stark contrast to the previous model, demand will always sum to potential output (and the interest rate will always equal its natural rate.) According to the model, prices are flexible so all prices adjust to clear the goods market; hence demand always equals supply. Yet how the composition of output changes over the business cycle *is* important. To give an example, suppose there is a once-off rise in A, but A will revert to trend next period. Clearly the rise in A causes output, Y_t, to rise. Now ask yourself, according to the permanent income hypothesis, what will happen? Because the change is temporary, most of the output will be *saved.* As in the Solow model, people use these savings to build up the capital stock (i.e., investment.) So what will happen is consumption will rise a little, but investment will rise a lot. For this reason, the model predicts moderately procyclical consumption, but highly procyclical and volatile investment. Moreover, because investment rises this period, the capital stock will be higher *next period*.[1] This then will lead to more consumption and more investment next period, and so on. For this reason, economic fluctuations will be *persistent*, as they are in reality.

There is more. As we shall see, the wage will equal the marginal product of labour, $\frac{\partial Y}{\partial L}$. But this is increasing in A_t. As a result, the temporary rise in A will also lead to a rise in the marginal product of labour: in turn, this raises labour demand and the real wage. This induces the *intertemporal substitution of labour*, causing labour supply to increase. So from

[1]Recall that $K_{t+1} = I_t + (1-\delta)K_t$. In the first period, the capital stock is *predetermined*. Investment *this* period raises the capital stock *next* period.

the production function, Y gets another "kick" from this increase in L.

Another important feature of the model is the way changes in A are temporary but modestly persistent. Say A usually grows at a rate of 2%. Then a "technology shock" would (say) cause A to grow at 4% in year one, 3% in year two, but would revert back to 2% in year three. Why do shocks have to be persistent? See, if shocks were just completely once-off, then almost all of the extra income would be saved when the shock hit. Remember the consumer is maximizing lifetime utility, so a temporary shock would only increase lifetime wealth—which, by the PIH, is what determines consumption in each period—by a little. To generate fairly procyclical consumption, as in the data, lifetime wealth must increase modestly. This is one reason why the shock to A must be persistent; namely, it must make the consumer feel moderately richer, but not *that* much richer. Yet this is not simply a modelling device: it is plausible that any change in A would last a few years. Whether the change is from innovation or a government policy, it seems reasonable that shocks to A would exhibit some persistence.

Implicit in the theory is the idea of strong and weak income and substitution effects. To induce a rise in both labour supply and saving today, the income effects associated with the technology shocks must be small. In contrast, the substitution effects must be *large*; this is what prompts the consumer to "make hay while the sun shines." By analogy with savings, if the wage rose permanently, then labour supply could in fact *fall* as a result of the technology shock (i.e., if the income effect was sufficiently strong). Again, the temporary nature of the shocks attenuates the strength of the income effects and reinforces the substitution effects.

Particularly important are the two key propagation mechanisms in the model: the *intertemporal substitution of labour* and *capital accumulation*. What I mean by a propagation mechanism is the model's internal way of amplifying the shock. In the New Keynesian model, sticky prices and real rigidity ultimately amplified shocks—and hence acted as a *propagation mechanism*.[2] Analogously, there are two key mechanisms here. First, consumers increase labour supply in response to a rise in the real wage. Because the technology shock also raises the marginal product of capital—and hence the natural rate of interest—consumers will also worker since they can now earn a greater return by purchasing capital and renting it out *next period*.[3] Second, because investment today leads to more capital

[2]The so-called *financial accelerator* acts as another propagation mechanism in the New Keynesian literature (though not in the simple version we studied.) Namely, increases in asset values in booms—as a result of lower interest rates—led to increases in consumers' net worth, which then induced banks to increase lending. This imparted a "multiplier effect" to any fall in interest rates. For firms that were borrowing, this also attenuated the rise in firms' marginal costs in booms. For this reason, the financial accelerator increased the degree of real rigidity in the economy, and lead to more powerful real effects of monetary policy.

[3]As noted, in any *given* period, the capital stock is predetermined.

tomorrow, it causes higher output again tomorrow. In turn, this leads to more investment tomorrow, so capital accumulation *propagates* the shock. In addition, as we know from last year, a greater capital stock raises the real wage, so this was induce more labour supply next period.

6.2 The Model

The representative consumer (or household) maximizes

$$E_0 \sum_{t=0}^{t=\infty} \beta^t \left(\log C_t - \frac{l_t^{1+\sigma}}{1+\sigma} \right)$$

where $\beta = \frac{1}{1+\rho}$, and ρ is *the rate of time preference*. The flow budget constraints each period are

$$w_t l + r_t k_t = C_t + i_t$$

$$k_{t+1} = (1-\delta)k_t + i_t$$

Everything above is in real terms. I have implicitly normalized the price level to one. Here, all saving is achieved through accumulating capital. The consumer receives income from renting out capital and from supplying labour. In any period t, the level of capital is predetermined. Investment today changes the capital stock *next* period. When the consumer saves, he therefore considers the rental rate on capital *next* period, not this period. In addition, it is common to explicitly impose a time constraint such as

$$l_t + h_t = 1,$$

where l_t refers to labour supplied in period t, and h_t (by a process of elimination) refers to leisure. Here, l refers to labour *hours*, not bodies.[4] These must add up to available time, which I assume is simply 1. Taking the the expected paths of wages and interest rates as given, the consumer maximizes lifetime utility.[5] Ignore the expectation sign for now; this just arises from the fact that A is uncertain in the future, which causes uncertainty about all

[4] *Realistically*, though, there is more than one person in any economy, so the proper measure of labour supply is lN, where l denotes labour hours per person, and N the number in employment. To derive the theoretical results, I will assume only a single person, but keep in mind that lN is the true variable that represents labour supply at the aggregate level.

[5] This just means that when a worker increases labour supply, he doesn't internalize the fact that the increase in supply will tend to lower the real wage at the national level.

future variables such as C, w, and r (which, as we shall see, will be functions of A.) Because this uncertainty has no qualitative effects on the dynamics, I will mostly ignore in deriving the optimality conditions.

Combining both constraints above gives

$$w_t l_t + r_t k_t = C_t + k_{t+1} - (1-\delta)k_t$$

$$w_t l_t + r_t k_t = C_t + k_{t+1} - k_t + \delta k_t$$

$$w_t l_t + (1 + r_t - \delta)k_t = C_t + k_{t+1}$$

$$\underbrace{w_t l_t + (1 + r_t - \delta)k_t}_{sources} = \underbrace{C_t + k_{t+1}}_{destinations}$$

Ignoring the expectations operator, the Lagrangian is then

$$L = \sum_{t=0}^{t=\infty} \beta^t \left(\log C_t - \frac{l_t^{1+\sigma}}{1+\sigma} \right) + \sum_{t=0}^{t=\infty} \lambda_t \left(w_t l_t + (1+r_t)k_t - C_t - k_{t+1} \right)$$

Maximizing w.r.t C_t gives

$$\beta^t \frac{1}{C_t} - \lambda_t = 0 \Rightarrow \beta^t \frac{1}{C_t} = \lambda_t. \tag{6.1}$$

Maximizing w.r.t C_{t+1}

$$\beta^{t+1} \frac{1}{C_{t+1}} - \lambda_{t+1} = 0 \Rightarrow \beta^{t+1} \frac{1}{C_{t+1}} = \lambda_{t+1} \tag{6.2}$$

Maximizing w.r.t l_t gives

$$-\beta^t l_t^\sigma + \lambda_t w_t = 0 \Rightarrow \beta^t l_t^\sigma = \lambda_t w_t \tag{6.3}$$

Maximizing w.r.t l_{t+1} gives

$$-\beta^{t+1} l_{t+1}^\sigma + \lambda_{t+1} w_{t+1} = 0 \Rightarrow \beta^{t+1} l_{t+1}^\sigma = \lambda_{t+1} w_{t+1} = 0 \tag{6.4}$$

Maximizing w.r.t k_{t+1} gives

$$-\lambda_t + \lambda_{t+1}(1 + r_{t+1}) = 0 \tag{6.5}$$

Finally, we have the usual *transversality condition*:

$$\lim_{t\to\infty} E_0 \beta^t u'(C_t) k_t = 0$$

Basically, if we value consumption at "the end" (i.e., $\lim_{t\to\infty} u'(C_t) > 0$), then we shouldn't leave capital left over. Instead, we should eat it.

Solution under Certainty

It is convenient to to set $\delta = 0$, so you can think of r_t as the real return *net of depreciation*. Combining (6.1), (6.2), and (6.5) gives the Euler equation

$$\frac{1}{C_t} = E_t \beta(1 + r_{t+1}) \frac{1}{C_{t+1}}.$$

Combined with the budget constraint, this implicitly gives C_t and savings in each period (which in turn will determine the consumer's level of investment).[6] Temporary rises in interest rates will cause consumption to fall and savings to rise, leading to more investment. This implicitly gives the household's optimal consumption path, and hence savings. Note that the interest rate that appears here is r_{t+1}. In this model, the household earns income by purchasing and then renting out the capital *next* period. But if the rental rate rises *next* period, the household will respond to that *today* by reducing consumption and investing. Keep in mind, then, that the Euler equation above gives us information about savings and *investment*. For a given level of income, if consumption falls, for example, then investment must rise.

Combining (6.1) and (6.3) gives the labour/leisure optimality condition

$$w_t u(C_t) = v'(l_t)$$

Or more explicitly

$$\frac{w_t}{C_t} = l_t^{\sigma} \tag{6.6}$$

Implicitly, this gives labour supply (or leisure demand, $h_t = 1 - l_t$.)

[6] Note that this is just the Euler equation

$$u'(C_t) = E_t \beta(1 + r_{t+1}) u'(C_{t+1}),$$

where $u(C_t) = \log C_t$

6.2.1 Intertemporal Substitution of Labour

Iterating forward (6.6) (or combining (6.2) and (6.4)) gives

$$\frac{w_{t+1}}{C_{t+1}} = l_{t+1}^{\sigma}$$

Then dividing this by

$$\frac{w_t}{C_t} = l_t^{\sigma}$$

gives

$$\frac{l_{t+1}^{\sigma}}{l_t^{\sigma}} = \frac{w_{t+1}}{w_t}\frac{C_t}{C_{t+1}} \tag{6.7}$$

Ignoring uncertainty, the Euler equation is

$$\frac{1}{C_t} = \beta(1+r_{t+1})\frac{1}{C_{t+1}}.$$

To see the idea here, set $\beta(1+r_{t+1}) = 1$. Then (6.7) becomes

$$\frac{l_{t+1}^{\sigma}}{l_t^{\sigma}} = \frac{w_{t+1}}{w_t},$$

and finally we get

$$\frac{l_{t+1}}{l_t} = \left(\frac{w_{t+1}}{w_t}\right)^{\frac{1}{\sigma}}.$$

This has a nice interpretation. Before going on, note that the consumer likes to smooth labour over time (as with consumption) due to increasing marginal disutility. But just as the consumer is "seduced" into deviating from the optimal consumption path via interest rate changes, the consumer will also deviate from smooth labour supply in response to fluctuations in wage changes. In other words, this is just like an Euler equation for labour. And just as the θ parameter mediated the response of consumption to interest rate changes, the σ plays a similar role here. What it means is the consumer will spread labour over time in response to changes in wages. For example, if the wage increased today relative to tomorrow, the consumer would increase labour supply today relative to tomorrow. But if the wage increased proportionally in *both* periods, there would be no change in relative labour supplies between both periods. In the RBC model, this condition is important. Namely, when A increases temporarily, the wage will also increase. And because the increase is temporary, labour supply will increase this period relative to the next one (or more generally, future ones). Observe too that a high σ will attenuate the degree of *intertemporal substitution of labour.* To replicate the dynamics of the business cycle, σ must to below in this theory.

Intertemporal Substitution and Interest Rates

So far, I have set $\beta(1+r_{t+1}) = 1$. But more generally, when $\beta(1+r_{t+1}) \neq 1$, we get

$$\frac{l_{t+1}}{l_t} = \left(\frac{w_{t+1}}{w_t} \frac{1}{\beta(1+r_{t+1})} \right)^{\frac{1}{\sigma}}$$

In this case, the rise in the interest rate in period $t+1$ will raise labour supply today. Think of it this way. A rise in the rental rate means there are more profits to be made from renting out capital next period. In response to this, the consumer should increase investment *this* period. And one means of doing this is to earn more income *today* (by working more) and using that income for investment *today*. Another way of saying this is the higher rental rate next period raises the return to working *this* period; in other words, it acts like an increase in the wage.

Labour Supply Levels

Let's talk about the *level* of labour supply for a moment. For now, ignore interest rates. When $u(C_t) = \log C_t$, the first order condition for labour reduces to:

$$\frac{w_t}{C_t} = l_t^{\sigma} \qquad \Rightarrow l_t = \left(\frac{w_t}{C_t} \right)^{\frac{1}{\sigma}}.$$

This implicitly defines the household's labour supply curve. As we saw a moment ago, if the wage rises today relative tomorrow, the household will supply more labour today *relative* to tomorrow. But what about the actual *levels* of labour supply? With a transitory change in the wage in period *t only*, there are two effects. First of all, the wage, w_t rises. Second, C_t will only rise a little, and certainly less than proportionally to the rise in w_t: because of the permanent income hypothesis, the consumer will smooth that once-off increase to the wage over his lifetime. As a result of both forces, the condition above implies that l_t will rise. For this reason, we can treat the relationship as an upwardly sloping relationship between l_t and w_t; i.e., the household's labour supply curve. To see what will happen in period $t+1$, consider

$$l_{t+1} = \left(\frac{w_{t+1}}{C_{t+1}}\right)^{\frac{1}{\sigma}}$$

Well, w_{t+1} didn't change. But due to the increase in w_t (the substitution effect) and consumption smoothing (the income effect), C_{t+1} will surely rise a little. As a result, l_{t+1} will fall by a tiny bit. Overall, therefore, l_t will rise a lot and l_{t+1} will fall a little, and so $\frac{l_{t+1}}{l_t}$ will certainly fall, as we predicted. (In fact, given the consumer is infinitely lived—and so a once-off wage increase in period t will only have a negligible effect on consumption each period—we often simply assume that l_t rises and l_{t+1} *stays the same* in response to a transitory increase in the wage in period t.) Yet looking at levels—$l_t = \left(\frac{w_t}{C_t}\right)^{\frac{1}{\sigma}}$—a permanent change in the wage will also cause C to rise permanently (both C and w will rise proportionately) and so will have no effect on levels of labour supply.[7]

Timing

Next, I turn to timing. According to RBC theory, the real wage rises in period t due to a technology shock, but stays higher than normal for a few periods. Because the technology shocks are persistent, A—and hence the wage—stays above trend for a few periods. In addition, the increase in A induces an increase in savings and capital accumulation. As we will see, this also causes the wage to rise in subsequent periods. For both reasons, once a technology shock strikes, the wage will remain above trend for a few periods. For example, the time series for the wage could be $\ldots, 10, 10, 100, 70, 50, 10, 10, 10, 10, \ldots$. In this case, the labour supply will jump up in period 3 and will remain above trend until period 6. Naturally, the level of labour supply in period 3 will be the greatest. Labour supply will revert to its long-run trend in period 6.

6.3 The Firm

Firms hire labour and rent capital from the household. To determine equilibrium wages and rental (interest) rates, we must look at the firm's optimization problem. In contrast to the New Keynesian model, the firm here is *perfectly competitive* and does not choose prices. As in perfect competition, the firm is a price taker. I normalize the constant price level to 1. All the firm does is choose the optimal combination of labour, L, and capital, K, in the production process.

The production function is Cobb-Douglas:

[7] As we noted in a problem set, a permanent rise in A would initially cause r_{t+1} to rise above its steady state, and this would cause a rise in labour supply initially, and the economy *transitions* to steady state.

$$Y = AK^{\alpha}L^{1-\alpha}.$$

Ignoring time subscripts, the firm's profit is

$$\pi = AK^{\alpha}L^{1-\alpha} - wL - rK.$$

The firm has revenues of $AK^{\alpha}L^{1-\alpha}$ and costs of $-wL - rK$. As already noted, I have normalized the price to one, so w and r refer to the real wage and real rental rate, respectively. The firm's revenues are then simply price times quantity produced, which is just equal to quantity produced, $AK^{\alpha}L^{1-\alpha}$. Note that the levels of capital K and labour L the firm will hire do not necessarily equal the level of capital and labour supplied by the household (k and l respectively.)[8]

Labour Demand

$$\frac{\partial \pi}{\partial L} = (1-\alpha)AK^{\alpha}L^{-\alpha} - w = MPL - w = 0$$

$$\Rightarrow w = (1-\alpha)AK^{\alpha}L^{-\alpha} = MPL \tag{6.8}$$

As always, the first order condition gives the optimal rule for the firm—in this case the optimal hiring rule. This means that the firm will hire labour up until the wage equals the worker's marginal product of labour. If the marginal product exceeded the wage, the firm would hire more workers (and vice versa). Manipulating the condition above gives an explicit equation for the firm's optimal level of labour demand, L^d:

$$L^d = \left(\frac{(1-\alpha)AK^{\alpha}}{w} \right)^{\frac{1}{\alpha}}.$$

As we'd expect, labour demand is decreasing in the real wage—so for a given K and A, this defines a downwardly sloping labour demand curve. Importantly, labour demand is *increasing* in the level of A and the level of K; these variables will shift the labour demand curve. Increases in these variables make workers more productive, and hence raise the attractiveness of hiring more people. Especially important for RBC theory is that a technology shock—a sudden jump in A—will *raise* labour demand.

[8] Yet this will be true in equilibrium.

Capital Demand

To obtain the optimal level of capital hired, the firm solves:

$$\frac{\partial \pi}{\partial K} = \alpha A K^{\alpha-1} L^{1-\alpha} - r = MPK - r = 0$$

$$\Rightarrow r = \alpha A K^{\alpha-1} L^{1-\alpha} = MPK \tag{6.9}$$

This optimal rule dictates that the firm should hire capital until the marginal product of capital equals the rental rate the firm faces. Manipulating the conditions above gives capital demand, K^d:

$$K^d = \left(\frac{\alpha A L^{1-\alpha}}{r}\right)^{\frac{1}{1-\alpha}}$$

A rise in A will raise the level of capital demanded. Namely, A makes capital more productive (computers, say, are more powerful now), making firms want more capital. An increase in labour, L, will also increase capital demanded. With more workers, each machine now becomes more useful—making firms want more of them.

6.3.1 General Equilibrium

As shown, labour and capital demand come from *firm*. Labour and capital supply comes from the household. To obtain market clearing (or general equilibrium) prices, we combine both sides of the market. Keep in mind that, in their decisions, the household and the firm take the wage and price as *given*. So the household never says: when I supply more labour, I'll push down the wage. If you think about it, this is realistic; everyone is such a small component of the market, they take prices as given.

To obtain the equilibrium wage, we combine an upwardly sloping household labour supply curve and the downwardly sloping labour demand curve. (Note that the households labour supply curve will be upward sloping: with *temporary* changes in wages—in this RBC model—higher wages will increase labour supply.) Formally, we will have

$$l_t = L_t^d,$$

or

$$l_t = \left(\frac{w_t}{C_t}\right)^{\frac{1}{\sigma}} = L^d = \left(\frac{(1-\alpha)AK^{\alpha}}{w}\right)^{\frac{1}{\alpha}}$$

The interaction of both relationships will give the equilibrium wage. To derive the equilibrium rental rate, we must note that in any given period, the household's supply of capital is *fixed*. Thus to get the equilibrium rental rate, I combine the household's inelastically supplied capital supply (a vertical line) with the firms downwardly sloping demand curve. In equilibrium

$$k_t = K_t^d.$$

These relationships will give us the equilibrium levels of capital, k and labour, l, together with the market clearing wage rate, w^*, and rental rate, r^*. In particular, equilibrium production will be

$$Y = Ak^{\alpha}l^{1-\alpha}. \tag{6.10}$$

From (6.9), we have the equilibrium rental rate

$$r^* = \alpha A k^{\alpha-1} l^{1-\alpha} = MPK \tag{6.11}$$

Keep in mind that we often refer to the rental rate as the interest rate. And from (6.8), we have the equilibrium wage

$$w^* = (1-\alpha)Ak^{\alpha}l^{-\alpha} = MPL. \tag{6.12}$$

Finally, it is easily shown that

$$\Rightarrow w^*l + r^*k = Y = Ak^{\alpha}l^{1-\alpha} \tag{6.13}$$

that is, the payments to the factors of production exhaust output.

6.4 Long-Run Equilibrium

Assuming no steady state growth, in the long run, the labour supply of each person is constant at

$$l_t = \left(\frac{w_t}{C_t}\right)^{\frac{1}{\sigma}}.$$

If there was sustained growth in A, then this would lead to equal growth of C and w, in which case, l would still be constant. Because labour hours have been *approximately* constant in recent decades, the model is therefore consistent with long-run trends.

To find the equilibrium capital ratio, look back at the Euler equation:

$$\frac{1}{C_t} = \beta(1 + r_{t+1})\frac{1}{C_{t+1}}.$$

Assuming no growth, in a steady state $C_t = C_{t+1}$, which implies that $\beta(1 + r_{t+1}) = 1$. The rental rate is also constant in steady state and from (6.11) equals $r^* = \alpha A k^{\alpha-1} l^{1-\alpha}$. Thus we have

$$\beta(1 + r_{t+1}) = 1 \Rightarrow \beta(1 + \alpha A k^{\alpha-1} l^{1-\alpha}) = 1 \Rightarrow \frac{k}{l} = \left(\frac{\alpha A}{\frac{1}{\beta} - 1}\right)^{\frac{1}{1-\alpha}}.$$

6.5 RBC: Review of Dynamics

What happens when there is a technology shock? The following outlines the central mechanisms behind RBC theory:

- Take, for example, a temporary reduction in the level of regulation. Because firms now spend less time filling out forms etc, there is now more produced for any given K and L. As a result, total factor productivity, A, jumps up. Having rational expectations, the household and firm know this change is temporary and somewhat persistent.

- Because this raises workers' marginal product, labour demand rises, which pushes up the real wage. In response, households increase labour supply. (Remember, the emphasis in this model is *supply*, not demand.) Namely, they know this wage increase is temporary, so the substitution effect dominates. This increase in labour supply and the increase in A raises the marginal product of capital, which raises the demand for capital, pushing up the rental rate. Additionally, the expected interest rate next period rises—the changes are persistent—next period, which also raises labour supply *this* period.

- Household income rises. Because the change is temporary, most of the increase in income is saved and used for purchasing more capital. This rise in the future expected interest rate induces the rise in savings. Yet, since the change is persistent, consumption rises a little. Consider the usual national accounts equation, $Y = C + I$. Because Y has increased a lot—due to A and L rising—while C has only risen a little, I (or savings) therefore increases a lot. Households smooth consumption intertemporally by building up their capital stock.

- The rise in I leads to more capital in the second period. Because of the stochastic process for A, falls a little, but is still above trend. Meanwhile, the capital stock is now higher due to last period's investment boom. For both reasons, the wage is still above trend. As a result, labour supply remains above trend.
- Over time, A reverts to trend, and the shock "dies out." The additions to the capital stock will also depreciate, and capital will revert back to its "normal" level.[9] All the mechanisms above also die out, and the economy reverts to trend.

6.6 Discussion

6.6.1 Empirical Evidence

Because $Y = AK^{\alpha}L^{1-\alpha}$, we can calculate growth in A from time series data on output, labour, and capital. We say A is the *Solow residual* from this decomposition. Empirically, growth in A is indeed associated with large fluctuations in Y. This is *a priori* evidence in favour of the RBC model, which claims *exogenous* changes in A lead to fluctuations in Y. Moreover, labour productivity, $\frac{Y}{L}$ is

$$\frac{Y}{L} = \frac{AK^{\alpha}L^{1-\alpha}}{L} = \frac{AK^{\alpha}}{L^{\alpha}}$$

According to RBC theory, therefore, labour productivity—as a result of a rising A—can indeed be procyclical as in the data. By contrast, the New Keynesian model predicts labour productivity is countercyclical: in that model, L rises, causing $\frac{Y}{L}$ to fall.[10]

However, New Keynesians counter this by saying that A is endogenous to the cycle. For example, if there were increasing returns, A would rise endogenously in a boom. In such a case, the boom would *cause* the change in A. (This is the usual reverse causation problem that plagues econometric work.) Take, for example, a shop like Subway Sandwiches. At lunch time, there's surely an increase in effort and production by the staff, as they make more sandwiches for office workers coming in for lunch. But would we say some "technology shock" is causing the increase in output? Hardly. Rather, we'd say the increase in

[9] For the higher level of capital to be maintained, the savings rate would have to rise permanently; however, because of the temporary nature of the shock, nothing happens to induce a permanent increase in savings.

[10] In the standard model, there is diminishing returns to labour, causing labour productivity to fall as L rises. In our derivation of the model, we assumed constant returns to labour. In this case, labour productivity would be acyclical.

demand is *causing* the increase in output—and, as a result, demand fluctuations are causing productivity fluctuations. Output is rising since the staff are now working a lot harder. Specifically, New Keynesians say labour and capital *utilization rates* rise in booms. More formally, Keynesians claim the production function should really take the form

$$Y_t = A_t (u_t K_t)^{\alpha} (u_t L_t)^{1-\alpha},$$

where u_t refers to the level of utilization or "effort." For any given L and K, a rise in u will raise output. Namely, by increasing effort, a higher u would act just like an increase in bodies or hours. To see this, manipulate the above expression to yield

$$Y_t = (u_t A_t) K_t^{\alpha} L_t^{\alpha}.$$

Thus the Solow residual now is uA. So, with this extension, the Solow residual is not necessarily capturing total factor productivity A—it's capturing u as well. New Keynesians claim u is rising *endogenously* in booms, not A. This way, labour productivity, $\frac{Y}{L}$, can be procyclical, *without* exogenous changes in A.[11]

Stabilization Policy

According to the real business cycle model, stabilization policy is counterproductive. Recall that stabilization policy aims to reduce the volatility of the business cycle; the objective is to moderate both booms and recessions. By raising interest rates, say, the central bank reduces investment in a boom and raises investment in a recession. Now suppose the RBC theory is correct. That is, productivity is fundamentally higher in booms and lower in recessions. As an example, say each unit of new capital yields 5 units of output in a boom (since A is higher), but only 2 units of output in a recession. In addition, workers are more productive in booms and less productive in recessions. Yet by trying to reduce investment/employment in booms and trying to raise investment/employment in recessions, stabilization policy is transferring resources/production from productive periods to least productive periods. As a result, stabilization policy is *counterproductive* and reduces welfare. As an analogy, imagine you're in great form and full of energy; then you *should* work harder; when you're down and lazy, you *should* stay in bed. It's silly, isn't it, trying to *induce* you to work less when you're full of energy and to make you work harder when you're having a bad day? Analogously, the business cycle simply represents optimal re-

[11] Of course, RBC'ers exact the ultimate revenge by saying Y causes the money supply, M, to rise endogenously in a boom—to be sure, a devastating blow to the Keynesian view that M is causing fluctuations in Y.

sponses by economic agents to changes in their economic environment—so things should be left as they are. In other words, what is, *is* efficient.

Although RBC theory claims money is always neutral, they do, however, agree that fiscal policy *can* raise output. But how it does so is not via the usual Keynesian demand-side channels. Rather, changes in government expenditure expand output by raising labour supply. New Keynesians claim government expenditure increases aggregate demand and hence production (since output is demand-determined). But according to RBC theory, output is always fixed at potential. What happens, they claim, is a rise in government expenditure makes people feel poorer: because of Ricardian Equivalence, people realize that *they'll* have to now pay a higher tax bill in the future. As a result, their lifetime wealth falls. And because households now feel poorer, they reduce consumption and raise labour supply: the negative income effect induces them to consume fewer consumption goods and less leisure. The corresponding increase in labour supply causes an increase in output and hence an increase in economic activity. Thus the mechanism by which output increases is completely different to the channels in the New Keynesian model. Yet again, the RBC theory always stresses incentives to supply labour and investment.

Chapter 7

The Ramsey Model

The Ramsey-Cass-Coopmans (or simply the Ramsey model) is a long-run model, which is closely connected with the RBC model. As in RBC, output always equals potential, and the interest rate always equals the natural rate.[1] This is basically the RBC model with fixed labour and with technology growing steadily or simply fixed (so there's none of the fluctuations in A that are crucial to the RBC model). From now on, I will simply assume A is fixed. Furthermore, I will assume there is no depreciation, and labour supply is fixed at one (i.e., $l_t = 1$ for all t). When we strip out these features from the RBC model, we have a model with people deciding between saving/consuming. The model is then typically used to analyze the accumulation of capital over time, and how the economy reaches to a steady state where capital is ultimately fixed at some equilibrium level. As you might suspect, the model works quite like the Solow model; it describes the evolution of potential output over time. Instead of having a fixed savings rate, however, savings in the Ramsey model are endogenized via the Euler equation. Because $l = 1$ and A is constant, I don't put time subscripts on them.

To keep things simple, I'll use logarithmic utility as in the RBC model. Then the Euler equation is

$$\frac{1}{C_t} = \beta(1 + r_{t+1})\frac{1}{C_{t+1}}$$

Knowing that $\beta = \frac{1}{1+\rho}$, this reduces to

$$\frac{C_{t+1}}{C_t} = \frac{1 + r_{t+1}}{1 + \rho} \tag{7.1}$$

[1]This is because all prices are flexible.

As in the RBC model, the budget constraint is

$$w_t + r_t k_t = C_t + k_{t+1} - k_t$$

From (6.13) above, we have $w_t l_t + r_t k_t = Y_t$, and so (noting $l_t = 1$)

$$Y_t = C_t + k_{t+1} - k_t.$$

Manipulating this yields

$$k_{t+1} = k_t + Y_t - C_t \tag{7.2}$$

i.e., capital next period equals capital this period plus savings (investment).

Because this is long-run model, it analyzes the accumulation of capital over time. Crucial to the Ramsey model is the idea that what's not consumed is *invested.* The Euler equation, (7.1), dictates that consumption will grow if the interest rate is above ρ. Idea is, as long as the interest rate is sufficiently high, consumers will continue to save, leading to consumption growth over time. Yet once the interest rate hits ρ, consumers couldn't be bothered saving any more: here, the interest rate—which induces them to save—equals the rate of time preference —which induces them to consume *today*. At this point, both forces offset, and the show is over: consumers consume all income and from equation (7.2), this puts an end to capital accumulation. From then on, the capital stock remains fixed, $k_t = k_{t+1}$.

Now, to determine how the economy evolves, we must determine the interest rate that prevails at each point in time. Fortunately, from (6.11) we can get the equilibrium interest rate at each point in time: $r^*_{t+1} = \alpha A k_{t+1}^{\alpha-1} = MPK$. Substituting this into the Euler equation, (7.1) above gives

$$\frac{C_{t+1}}{C_t} = \frac{1 + \alpha A k_{t+1}^{\alpha-1}}{1 + \rho}$$

The story goes as follows. At the outset of development, the capital stock is low. And because of diminishing marginal product of capital, the marginal product of capital—and hence the rental/interest rate, r—is *high.* As a result, consumers save more. Yet as they save and accumulate more capital over time, the marginal product of capital and the interest rate, r, falls. And eventually the interest rate will fall to ρ. Once this happens, the consumer no long saves: the tug of war between interest rates and impatience is over. At this point, there is no way to increase potential output, so output growth stalls. This level of output becomes the steady state output. In steady state, the consumer then simply consumes all of some constant output level.[2]

In steady state, we know consumption will be constant when

$$\alpha A k_{t+1}^{\alpha-1} = \rho$$

$$\Rightarrow k^* = \left(\frac{\alpha A}{\rho}\right)^{\frac{1}{1-\alpha}}$$

Then from (6.10)—and noting $l = 1$—we can get equilibrium output, $Y = Ak^\alpha = A\left(\frac{\alpha A}{\rho}\right)^{\frac{\alpha}{1-\alpha}}$. Note that a higher rate of time preference—indicating more impatience—implies the level of the steady state capital stock is lower. By reducing the incentive to save, a high ρ leads to less capital in steady state.

Taxation and Ramsey Model

We now consider taxes. If we place a tax on capital, then the after-tax interest rate—which determines decisions—is $(1-\tau)\alpha A k_{t+1}^{\alpha-1}$, and so the steady state condition is

$$(1-\tau)\alpha A k_{t+1}^{\alpha-1} = \rho$$

The capital stock then becomes

$$\Rightarrow k^* = \left(\frac{(1-\tau)\alpha A}{\rho}\right)^{\frac{1}{1-\alpha}}$$

Why does the steady state level of capital fall? The increase in taxation reduces the after-tax interest rate, which reduces the incentive to save as the economy develops. In turn, this reduces capital accumulation and growth. And because wages are increasing in the level of capital—since capital raises the marginal product of labour—higher tax rates will also lead to lower wage growth along the path to steady state and in steady state. However, this after-tax interest rate remains fixed at ρ. The reduction in the capital stock raises the marginal product of capital, and this combined with the higher tax rate yields the same after-tax rate.

If an economy is in steady state, and taxes on interest rates fall, then the consumers will save and accumulate capital until we reach a new higher steady state capital level. In a more general setting—without labour fixed—lower tax rates in interest rates would also raise labour demand and employment.

[2] Of course, if A grew, standards of living would rise, just as in the Solow model.