Review Notes

Paul Scanlon

EC 4010

Contents

Chapter 1

The Ramsey Model

The Ramsey-Cass-Coopmans (or simply the Ramsey model) is a long-run model, which is closely connected to the real business cycle (RBC) model. Actually, it *is* the RBC model, where both the level of total factor productivity, A, and labour supply are fixed. As in the RBC model, output always equals potential, and the interest rate always equals the natural rate.[1] To keep the presentation simple, I assume there is no depreciation, and labour supply is fixed at one (i.e., $l_t = 1$ for all t); conveniently, this means everything will be in *per-capita* terms. After we have stripped out these features from the RBC model, we have a model where people simply decide between saving and consuming. Because people use savings to finance investment, this enables us to analyze the accumulation of capital over time and, in turn, economic development. Eventually the economy reaches a steady state where capital is fixed at some equilibrium level. In this sense the model works quite like the Solow model; it describes the evolution of capital and hence potential output over time. Yet instead of having a fixed savings rate, in the Ramsey model we endogenize savings via the Euler equation. But the insights are the same.

From now on, I'll use logarithmic utility as in the RBC model. So the Euler equation is

[1] This is because all prices are flexible.

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

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

$$\frac{C_{t+1}}{C_t} = \frac{1 + r_{t+1}}{1 + \rho} \tag{1.1}$$

As in the RBC model, the budget constraint is

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

At the outset, the initial level of capital, k_0 is given. In the RBC model, payments to the factors of production exhaust output, so $w_t l_t + r_t k_t = Y_t$. Combining this with the constraint above and noting $l_t = 1$, we have

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

Manipulating this yields

$$k_{t+1} = k_t + Y_t - C_t \tag{1.2}$$

i.e., capital next period equals capital this period plus savings, $Y - C$ (i.e., investment).

Seeing that this is a long-run model, our main focus is the accumulation of capital over time. The Euler equation, (1.1), dictates that consumption will grow if the interest rate is above the rate of time preference 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 stop saving. At this point, the show is over: consumers consume all income and from equation (1.2), this puts an end to capital accumulation. From then on, the capital stock remains fixed, $k_t = k_{t+1}$. At this point, the economy is in steady state.

Now, to determine how the economy evolves, we must determine the interest rate that prevails at each point in time. Fortunately, from the RBC analysis, we can get the equilibrium interest rate at each point in time: $r^*_{t+1} = \alpha A k^{\alpha-1}_{t+1} = MPK$. Substituting this into the Euler equation, (1.1) above gives

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

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, causing consumption to grow. 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 the interest rate r and impatience ρ is over. Considering that there is no other way to increase potential output—recall that labour and TFP are fixed—output growth stalls. This level of output becomes the steady state output.[2] In steady state, the consumer then simply consumes all of some constant output level.

In steady state, we know consumption will be constant when

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

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

Whatever happens, the steady state interest rate always equals ρ. Now, from the production function—and noting $l = 1$—we derive 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 at each point in time, a high ρ leads to less capital—and hence output—in steady state.

[2] A more detailed analysis would demonstrate that the economy in fact settles here in equilibrium. For instance, you could argue that agents start "eating" the capital stock at this point, causing negative growth. A more rigorous analysis rules out such a possibility.

1.1 Other Issues/Applications

Taxation and Ramsey Model

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 seeing that 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.

If an economy is in steady state, and taxes on interest rates fall, then consumers will save and accumulate capital until the economy attains a new higher steady state capital level.

The Small Open Economy

Consider two economies, a small one and a large one. Recall first that the expression for the rental rate for capital in any economy is given by

$$r = \alpha A k^{\alpha-1}. \tag{1.3}$$

Expressed this way, the rental rate for capital is increasing in A and *decreasing* in capital per person. It follows that if a country has a low level of capital per person, r is higher. However—and this is important—if the country *also* has a low level of

A, that would reduce r again. Accounting for this adjustment, we see that countries with low levels of capital per person do not necessarily offer high returns.

Now suppose the average interest rate in the rest of the world is $\overline{r}$ and the rates of time preference ρ and A are the same everywhere. If $r > \overline{r}$, arbitrage ensures that capital will flow in to the small open economy; and if $r < \overline{r}$, capital will flow out. That is, a low domestic interest rate causes a *capital/savings outflow* as investors exploit better investing opportunities abroad, and vice versa. This will continue until the interest rate in the small open economy is equal to the "world" interest rate. So, in equilibrium with perfect capital mobility, the interest rate is taken as given, and we have

$$r = \overline{r},$$

where $\overline{r}$ is the *world* interest rate.
Combining this equilibrium condition with (1.3) above gives:

$$\overline{r} = \alpha A k^{\alpha - 1}.$$

In this case, the level of capital in the small economy would converge immediately to that of the large economy: interest (i.e., capital) difference would be arbitraged away. For a large open economy, it would still be reliant on its own national savings. Clearly in this case, there wouldn't be sufficient flows from abroad to make its interest rate *independent* of domestic fundamentals. Now, recall from above that if a country has a low A, capital will not flow in even if k is relatively low. The fact capital does not flow to small developing countries is further evidence that A is extremely low in the developing world.

Aside: Growth in TFP

Returning to the theme of growth, the model provides a good description of the earlier stages of development and predicts the kind of conditional convergence we see in the data. Nonetheless, it fails to predict sustained increases in living standards.

Certainly this is a major flaw for a growth model. Like in the Solow model, we address this by assuming long-run growth is driven by exogenous growth in total factor productivity A. I will not formally incorporate this into the model, but, sure enough, it can be easily done, with the insights above and below remaining intact. As we shall see, the Romer model of endogenous growth addresses this important issue.

Chapter 2

Endogenous Growth: The AK Model

What I mean by endogenous growth is long-run growth determined by economic factors within a model. To generate long-run growth in the Solow/Ramsey model, we simply *assumed* A grew over time; there, long-run growth was exogenous. But, in this sense, the model's implications for sustained growth were unsatisfactory: the model explained everything except the reason for long-run growth itself. Now I turn to the AK model. Although the empirical support for this model is weak, it does generate growth endogenously, and provides a useful introduction to the endogenous growth literature.

To begin, the aggregate production function in the Ak model takes the *linear* form:

$$Y = Ak, \qquad A > 0,$$

where A is *constant*. You might ask, well where does the Ak model come from? One obvious and uninteresting way to generate the $Y = Ak$ formulation is simply to consider the model a special case of the Ramsey model. There, production was given by $Y = Ak^{\alpha}L^{1-\alpha}$, so by imposing $\alpha = 1$—i.e., no diminishing returns to capital—we get the Ak model. Still, you might argue that this seems odd: the argument that a tenth laptop, for example, is as useful as the first one seems untenable. So next I turn to a more realistic way to generate a production function $Y = Ak$.

Learning By Doing

Suppose production is given by:

$$Y = Bk^{\alpha}, \qquad \alpha < 1,$$

where B denotes productivity. For convenience, we assume labour is of no use in production. The central twist here is that productivity is a function of capital:

$$B = Ak^{1-\alpha}, \qquad A > 0,$$

that is, B is increasing in the level of the capital stock. Point is, as workers work with more capital, they become more skilled; this is what we call "learning-by-doing." (Imagine your mind becoming "quicker" and more agile as you work with more equipment.) Then substituting the expression for B into the production function gives

$$Y = \underbrace{Ak^{1-\alpha}}_{\text{externality}} k^{\alpha} = Ak,$$

which is our friend, the Ak model. It is instructive to consider how this works. Raising k has two effects on the marginal product of capital. First, a rise in k has the usual diminishing returns effect of *reducing* the marginal product of capital; this is evident from the production function, since $\alpha < 1$. Second, a rise in k raises productivity, B; this *raises* the marginal product of capital. And, overall, both effects just offset, yielding constant returns to capital. On net, therefore, the rising productivity "kills" the diminishing returns effect.

Technically, we say that there is a positive externality to physical capital accumulation. As we will see with the Romer model, this kind of feature leads to a social inefficiency: individual firms don't internalize the productivity gain to their individual investments and, as a result, they will underinvest in capital. But more on this later.

Another way to generate an Ak model is to incorporate human capital. In this formulation, there are two kinds of capital: physical capital and human capital. The idea here is that increases in human capital raise the marginal product of physical capital. Imagine, for instance, firms getting more "bang per buck" per unit of capital when there are more skilled workers working with the capital. And if levels of human capital rise over time along with physical capital, this can ultimately offset diminishing returns to physical capital. Ultimately, we end up with $Y = Ak$ model, where k is now interpreted broadly to encompass both physical *and* human capital.

2.0.1 Factor Prices

Recall now that the wage and interest rate are determined by the usefulness of labour and capital (respectively) in production; i.e., by their marginal products. Because raw labour is of no use in the $Y = Ak$ production function, the marginal product of labour, $\frac{\partial Y}{\partial L}$, and hence the wage is 0. More importantly, the marginal product of capital, $\frac{\partial Y}{\partial k}$, is A, which of course is constant. Because there's no diminishing returns to capital in the Ak model, the usefulness of capital—i.e., its marginal product—remains constant over time. As we shall see in a moment, this will change the dynamics of the model dramatically. Concerning income distribution, the total payments to the factors of production are

$$wL + rk = \frac{\partial Y}{\partial L}L + \frac{\partial Y}{\partial k}k = 0.L + Ak.$$

And since $Y = Ak$, reassuringly we have

$$wL + rk = Y,$$

that is, payments to the factors of production exhaust all of GDP. The owners of capital receive all the income in this Ak economy.

2.0.2 Dynamics

To analyze dynamics, assume for simplicity that there is no population growth. This lays bare the central mechanism of the model: capital accumulation. The per-capita production function is:

$$y = Ak,$$

and the Euler equation is

$$\frac{C_{t+1}}{C_t} = \frac{1+A}{1+\rho}$$

Because there are no diminishing returns, the interest rate remains the same, regardless of the amount of savings. For this reason, people will continue to save (and hence consumption will continue to grow.) Seeing that the capital stock will be growing forever, we therefore get perpetual growth. In this sense, long-run growth is endogenous to the model.

And because $y = Ak$, where A is a constant, y and k must grow at the same rate. Hence,

$$\frac{\dot{k}}{k} = \frac{\dot{y}}{y}. \tag{2.1}$$

That is, the rate of output growth equals the rate of capital growth. In this economy, income per capita *always* grows at the *same* rate. But more important, long-run growth is now *endogenous*—and this can occur even without growth in A. So, on the *balanced growth path*, income per person rises at a constant rate. (Of course, to match the reality that living standards rise over time, we must impose the condition $A > \rho$.) Significantly, and in contrast to the Ramsey model, the *growth rate* of income per capita—i.e., the standard of living—is now increasing in A. Why? Because there are no diminishing returns to capital accumulation, the return to savings—and hence investment—never declines in value. By contrast, in the Ramsey model, savings only had *level* effects as a result of diminishing returns to capital. Unhappily for us, the model fails to predict conditional convergence;

i.e., the fact that poor countries with good steady states grow faster. To attain this feature, a model must exhibit diminishing returns to capital (holding TFP constant.) For this reason, this model does not provide a satisfactory explanation for why economies grow over time. Nonetheless, we have learned something important: to generate long-run growth endogenously, a model must have some feature that "kills" diminishing returns to capital accumulation.

Chapter 3

Endogenous Growth: The Romer Model

Although the Ak model generates long-run growth endogenously, it fails to fit the data along important dimensions. Another more successful model of endogenous growth is the Romer model, developed by Paul Romer of Stanford. While the emphasis in the Solow and Ak models was factor accumulation, or *objects*, the emphasis in this model is on *knowledge* and innovation. This model and its offshoots—the *"New Growth Theory"*—were formulated to explain where technology or A came from; i.e., to *endogenize* technological progress.

Central to the model is the idea that firms introduce new capital goods over time. Especially in the developed world, this accords with reality; firms introduce new types of goods all the time and this does seem to play a central role in driving growth. One important implication of the model is that some *monopoly power* is necessary to stimulate purposeful innovation, and therefore sustain growth. That is, we must give people *incentives* to undertake research and improve technology. Romer was the first one to introduce the idea of *imperfect* competition into growth theory, and to emphasize the centrality of the profit motive to sustained growth.

First, I present the model at the economy-wide level; after that I present the microfoundations of the model. But a few caveats before we start. For our purposes, it's best to see Paul Romer taking up where the Ramsey model reaches steady state; so instead of being substitutes, think of these models as *complements* More important, we should view this as a model for the *developed* world, which should not be applied seriously to developing countries; after all, almost all technology is developed in advanced nations. Also, because knowledge and ideas can diffuse across borders, the appropriate unit of observation for the model is the world, not individual countries. Finally, the A in this model is quite specific: it refers to the number of new capital goods that have been invented. But more broadly, however, you can view A as the stock of useful ideas, knowledge, and innovations; all of those are effectively "recipes" for rearranging capital resources to generate more output.

3.0.3 The Model

In the Romer model, there are two sectors: the production sector and the research section. To see the main idea, consider the following example.

Imagine you are the only person involved in production—i.e., $L = 1$—and you are working with a piece of capital, k_1:

$$Y = k_1^\alpha, \qquad 0 < \alpha < 1,$$

which is clearly subject to diminishing returns. But then a new piece of capital, k_2 (of the same price), is introduced, so your production function becomes:

$$Y = k_1^\alpha + k_2^\alpha.$$

Now, ask yourself: how do you allocate expenditure across the capital goods? Well, if there's diminishing marginal product to each good, the best thing is to spend an equal amount on each good; this way, you equate the "bang per buck" from each good. (Just think: this is reason you would rather have a laptop and a desktop, as opposed to two desktops.) More generally, rather than using one capital good, it's optimal to employ *all* existing ones.

If you think about it, this is quite similar to the way people smooth consumption over time. Here, at the aggregate level, the economy is smoothing capital resources over all existing kinds of capital. According the permanent income hypothesis, people smoothed *consumption* across time because of diminishing marginal utility to consumption *in each time period.* Analogously, here we are "capital smoothing" as a result of diminishing marginal product of capital *for each unit.* Really, it is the same idea.

Most importantly, when you start off with k_2, the marginal product of that new unit is extremely high (technically it is infinite, since $\frac{\partial Y}{\partial k_2} = \frac{\alpha}{k_2^{1-\alpha}}$, which is infinite when $k_2 = 0$). In these sense, the introduction of new forms of capital—such as computers, laptops, printers, and so on—can *offset* diminishing returns to capital. As such, one can view the basic Solow model is a special case of this model, where you are just working with one kind of capital good; and growth falters because no new capital goods are introduced.

The Complete Model: The Production Sector

Following on from the above, the aggregate production function in the Romer model is given by:

$$Y = L_y^{1-\alpha} \sum_{i=1}^{i=A_t} k_i^\alpha,$$

where L_y denotes the number of workers involved in production, and A_t denotes the number of capital goods available at time t. (Content yourself that the "baby" production function above was a special case of this where $A_t = 2$, and $L_y = 1$.) The trick here is that A increases over time. Firms continually introduce new varieties of capital; and all new capital goods are of equal value and treated symmetrically. Thus capital resources are continually being put to new uses.

Suppose now that capital resources are K; this is the level of real expenditure you have set aside to purchase capital goods. Then, by "capital-smoothing" we will allocate an equal amount of expenditure to each existing capital good. Hence

$k_i = \frac{K}{A}$, and so the level of production reduces to

$$\begin{aligned} Y &= L_y^{1-\alpha} \sum_{i=1}^{i=A} \left(\frac{K}{A}\right)^{\alpha} \\ &= L_y^{1-\alpha} \underbrace{\left(\left(\frac{K}{A}\right)^{\alpha} + \left(\frac{K}{A}\right)^{\alpha} + \ldots \right)}_{A \text{ of these}} \end{aligned}$$

And by symmetry:

$$\begin{aligned} Y &= L_y^{1-\alpha} A \left(\frac{K}{A}\right)^{\alpha} \\ &= A^{1-\alpha} K^{\alpha} L_y^{1-\alpha}. \end{aligned}$$

So the Romer production function reduces to one almost identical to that in the Ramsey model. The fact that A has an exponent of $1-\alpha$ and L_y replaces the usual L are the only differences, but these have no qualitative bearing on the results. So for convenience we can treat this production function as identical to the standard one, with one important exception: now A has a story behind it.

Because capital enters as K^{α}, notice that there *are* diminishing returns to capital *resources/expenditure*. But, crucially, the marginal product of capital is increasing in the level of A; i.e., A offsets diminishing returns. And as long as A is continually rising, this can offset diminishing returns indefinitely, and ensure sustained growth. Notice above that if we double K *and* double A, output will also *double*. That is, with A growing, the production function behaves *as if* there are no diminishing returns to capital. But, with A fixed, we are back to diminishing returns again. Continual growth in A is how the Romer model generates long-run growth.

The Research Sector

Having shown that A is important, it is natural to ask: where does A come from, anyway? And how do we know A will keep growing? In the model, increases in A come from the *research sector*, to which I now turn.

Suppose that the population is L, and there is no population growth. The population is divided into those employed in production, L_y (which we have already seen) and those employed in research, L_A:

$$L = L_y + L_A.$$

Assume further that a fraction $0 < v < 1$ of the population are involved in research so:

$$L_A = vL$$

and then by a process of elimination:

$$L_y = (1 - v)L.$$

Most important, there is a labour-intensive research sector, where the production function for innovations is

$$\dot{A} = \gamma A L_A = \gamma v A L.$$

Notice that the change in A is increasing in the level of A. Existing innovations help stimulate new research; the development of the laptop, for instance, was surely aided by the previous development of the desktop computer. In other words, there are positive externalities to the existing stock of A. To obtain the growth rate of A—the source of long-run growth—divide across by A to get

$$\frac{\dot{A}}{A} = \gamma v L,$$

which is constant. Notice that the growth of A—and hence the standard of living—is increasing in the level of the population. And over the course of history, there is a positive relationship between growth and population, lending support to this idea. So that's it—this is essentially the Romer model.

3.0.4 Discussion

The idea that long-run growth is increasing in the level of population is a novel one and demands further scrutiny. What is the source of this result? Underlying this result is the fundamental idea that innovations are *nonrival* goods. To say a good is nonrival means one's own use of it does not preclude someone else from using it. Consider, for example, a computer: the design for the computer—part of the stock of A—can be used freely by anyone, but a physical computer can only be used by one person (formally, we say a computer is a *rival* good). As another example, consider these notes. Your reading of them does not in any way preclude anyone else's. I can put them on the web and 6 billion people could read them. What matters for your education is not the number of books written per person, but rather the *stock* of books in print. To repeat, it is different for the computer: only *you* can use it.

More formally, what matters for welfare is capital per person, not the stock of aggregate capital. But because innovations can be shared by everyone, what matters is the *stock* of existing innovations, not innovations per person. Precisely because of this, the model has the implication that more people are better: more people lead to more innovations, which can then be *shared* by everyone. This feature is called a *scale effect*. To see this formally, consider the standard per capita production function:

$$y = \underbrace{A}_{\text{can be shared}} k^{\alpha}$$

In a world where technology diffuses (knowledge leaks), what matters for individual welfare is the aggregate stock of A. But for rival goods like capital, what matters is capital per person, k. Because K is relatively unimportant for long-run growth, while A *is* important, this is good news for developing countries. By importing A, there is the possibility to develop quickly.

So we now have a model of long-run growth, where growth is determined endogenously within the model. While the Romer model can explain long-run growth, the Ramsey model can explain the levels. Therefore, our complete model of growth will just be the Romer model appended to the Ramsey model. At the outset of development, this will behave just like the Ramsey model, exhibiting conditional convergence. Then, once in steady state, you can think of the Romer model then taking over, determining the dynamics on the balanced growth path. (The phrase "balanced growth path" just refers to the long-run dynamics of a model.)

Finally, you might wonder: considering $\frac{\dot{A}}{A} = \gamma v L$, do growth rates rise over time if the population level rises? Well, yes, according to this model. You might argue that this is at odds with the data, and it is. In case you are worried about this, don't be: one variation of the Romer model, which corrects this "defect" is to posit a knowledge production function of the form $\dot{A} = \gamma A^{\phi} L_A^{\lambda}$. Here ϕ mediates the extent of knowledge spillovers, while λ mediates the usefulness of researchers. Notice, in particular, that the Romer model is a special case of this where $\phi = 1$ and $\lambda = 1$. This more general formulation of the Romer model can generate the more reasonable result that long-run growth is increasing in the rate of population *growth*, as opposed to the population *level*.[1] But fortunately the essential insights of the basic Romer model—which I will stick with—remain intact.

3.1 The Economics of Innovation

What we have seen so far is a birds-eye view of the economy; we have examined the Romer model, without presenting any microfoundations. Now I turn to the microeconomics behind the model. To start with, recall what happened behind the scenes in the Solow model: there were firms maximizing profits, and individuals supplying labour and renting capital to firms. And, in that setting, firms were price takers. Now I discuss firms in the Romer model. In contrast with the Solow model, firms here must have some monopoly power. Because innovation entails a large fixed costs, firms must be assured profits to innovate. And this raises a number of issues about the economics of innovation. Further, the nonrivalrous nature of ideas makes the microeconomics of the model somewhat different: there

[1] The model is then called a semi-endogenous growth model.

are externalities to innovation, and, as we shall see, this will introduce economic inefficiency in equilibrium.

3.1.1 Fixed Costs and Pricing

The novel thing about innovations is that they entail a large fixed cost, but once made, entail small marginal cost. Mathematically, suppose our *fixed cost*—here, the cost to invent a new product—is F and the constant *marginal cost* is w. Our total cost function for producing x units of output is therefore:

$$C = F + wx.$$

Average cost, AC, is then:

$$AC = \frac{C}{x} = \frac{F}{x} + w.$$

Because of the fixed cost, average cost always exceeds marginal cost:

$$AC = \frac{F}{x} + w > w = MC.$$

For this reason, firms make losses with *marginal cost pricing*, since average cost would now exceed price. Now, here's the thing: the most celebrated model in micro is that of a perfectly competitive *price taking firm* taking the price of its product as a *given* and setting price (which equals marginal revenue under perfect competition) equal to the marginal cost. This maximizes total consumer and producer surplus, and entails no deadweight loss. But not so here. Contrary to what happens under *perfect competition* (where $P = MC$), here we must have $P > MC$ to cover the fixed costs of innovation, F. In equilibrium, price equals average cost. Thus covering fixed costs requires *imperfect competition*—i.e., the price must be *greater* than marginal cost—and the firm must have some market power to set prices. This wedge between willingness to pay (i.e., demand) and the marginal cost of production gives rise to a deadweight loss and social inefficiency. In particular, we must have *monopoly rents* of $P - MC$ on all sales to cover those initial fixed costs. With perfect competition, there will be no innovation and no growth.

Therefore, when we discuss pricing decisions in the Romer model, we must model firms having market power. In his formal model, Romer modeled firms as monopolistic competitors, all of whom had some monopoly power.[2] But there is a problem: given ideas/recipes are nonrivalrous, how do we prevent other firms from selling your inventions? That is, how do we ensure the firm can keep its idea/innovation to itself, so it can recoup its fixed cost? Well, we must have intellectual property rights—such as patents and copyrights—that guarantee firms at least temporary monopoly power over their innovations. Thus the Romer model has the implication that intellectual property rights and some degree of monopoly power are necessary to sustain growth. In the case of innovation, we say there is static inefficiency, but dynamic efficiency: although there are distortions—due to monopoly power—at any point in time, this is necessary to encourage innovation and sustain growth (hence the dynamic efficiency part).

One important corollary of the above result is that the First Welfare Theorem—i.e., the result that the market produces the socially efficient, Pareto optimal outcome—doesn't hold. Why? The price exceeds the marginal, so the equilibrium cannot be Pareto optimal. To understand this, imagine that the price of good you like is 1000, but the marginal cost is only 1. Like most students, you can't afford 1000, but would be willing to pay 50 for the good. Although it is socially optimal for this transaction to occur—your willingness to pay, 50, exceeds the marginal cost of production, 1—this exchange does not take place. As a result, a situation like this results in a loss in consumer surplus—a pure deadweight loss. The point is, by lowering the price and increasing output, we can make you and the producer better off, and attain a Pareto superior outcome. Because price equals marginal cost under perfect competition, we don't have this problem. To conclude, the problem in the Romer model is that there is no socially optimal, "right" price. On the one hand,

---

[2]Monopolistic competition is a somewhat milder form of monopoly. While a monopoly has no competition, a monopolistic competitor does: there are other firms around selling similar, but not identical products. To take an example, all restaurants in a food court are monopolistic competitors. McDonalds and Burger King, for instance, both have market power, but not much.

we need a high price price to encourage innovation; on the other, economic efficiency dictates that socially optimal price is almost zero. So a single price is trying to do two very different things.

Result 1 *To spur innovation, firms must have the ability to set a price greater than marginal cost. Therefore, we need some form of imperfect competition; with perfect competition, there will be no innovation and no sustained growth. And because the price exceeds marginal cost, the equilibrium is not Pareto efficient.*

3.1.2 Externalities

Another market failure stems from the positive externalities associated with knowledge production. An important concept when discussing the nonrivalrousnss of ideas is the concept of *social* versus *private* return. The private return is *your* return, whilst the social return is the overall return and benefit to *society.* When a good betrays nonrivalrousness, the social return will typically be greater than the private return. And one important feature of the production function for ideas—$\dot{A} = \delta A L_A$—is that there are positive externalities to innovation; that is, the social marginal return to innovation exceeds the private marginal return. Yet private researchers will not internalize the complete consumer surplus associated with their innovations in their private decisions. Instead, they will just consider the private gain to innovation—their profit. As a result, they will underinvest—relative to the social optimum—in research and development. Consider, for example, the Solow model. When Robert Solow wrote his seminal paper on growth, I suspect his private return from the publication was about $50 or so. Yet, considering the joy that millions of students have received from his insights ever since, then we can see how the social return to his work was greater than the private return.

This is *another* reason the equilibrium in the Romer model is not Pareto optimal. Because there are positive externalities to research, a benevolent social planner would want to increase the number of researchers, L_A, in the economy. Concerning possible action in response to this market failure, one possible resolution is for the government to subsidize industries engaged in research. Ideally this would increase the number of researchers towards the *social optimum*.

Result 2 *Because there are positive externalities to innovation, the equilibrium level of innovation will be suboptimal and inefficiently low. For this reason, the equilibrium is not Pareto optimal: welfare in the economy could improve if there was more research and development.*

3.2 Knowledge Diffusion

Our combined Solow-Romer model predicts that developed countries—where innovations diffuse seamlessly—should have the same long-run growth trend, but different levels (depending on factors like national savings rates.) And consistent with the model, developed countries indeed share approximately the same trend growth, but have different income levels.

Chapter 4

Public Finance

4.0.1 Optimal Taxation

Because they induce people to alter their decisions in an attempt to avoid them, taxes are *distortionary.* One common example of a non-distortionary tax is a head-tax: everyone, say, must pay 100 to the government, regardless of income. Seeing that they minimize deviations from a Pareto optimal equilibrium, economists like such non-distortionary taxes. An important question is how to raise revenue while minimizing distortions in economic activity. The theory of *optimal taxation* addresses this issue.

Before going on, it is important to note that the costs of taxation are convex. That is, as tax rates rise, then the associated distortions increase disproportionately. To see this, look at Figure 1.1. We initially start off with supply equal demand. Suppose the government now introduces some tax on consumers who purchase the good: for every unit purchased, they must pay an addition t to the government. So if consumers were happy to purchase Q at price p before, they'll now only be happy to purchase Q at price $p-t$. This happens for every Q along the x-axis. As a result, the demand curve shifts down by the size of the tax, t. Now recall that the demand curve indicates willingness to pay, while the supply curve indicates marginal cost. We say that the triangle in the diagram is the *deadweight loss* associated with the tax. Namely, for all the quantities below the triangle, the willingness to pay (i.e., demand) exceeds the marginal cost (i.e., supply). All of these transactions took place initially, but not anymore. Therefore, the triangle represents the loss in social welfare associated with the tax. Continuing this analysis, take a look at Figure 4.2. If we raise taxes *again* by t we get another shift down. Most importantly, note that the size of the deadweight loss grows disproportionately—it *more* than doubles. Finally, when demand is relatively inelastic (as in Figure 4.3), the associated deadweight loss is smaller.

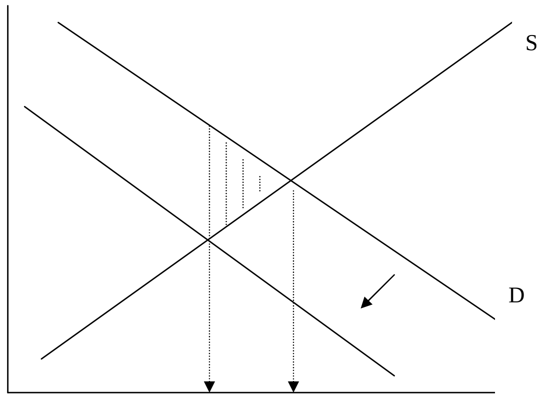

Figure 4.1:

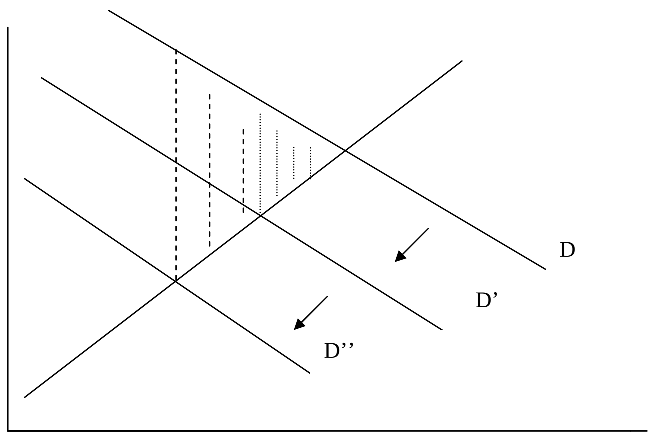

Figure 4.2:

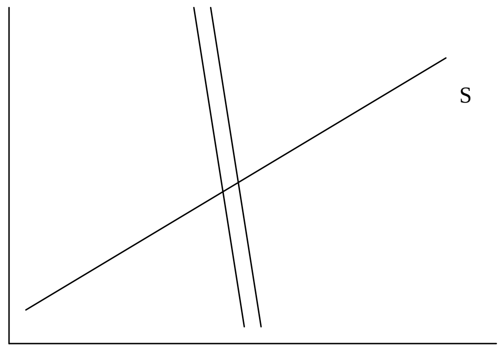

Figure 4.3:

These diagrams have a number of broad implications. First, the government should spread the burden of taxation as thinly as possible; it should keep rates low and broaden the tax base. To understand this, take a look at Figure 4.2: increasing taxes too much on a given activity causes disproportionately larger distortions. Second, to minimize distortions, the government should tax activities that are relatively inelastic. In Figure 4.3, a relatively inelastic demand curve results in a lower deadweight loss. This concept is often referred to as the *Ramsey rule*: tax goods that are inelastically supplied. By definition, such goods are relatively insensitive to price changes. By taxing such goods, we minimize distortions, and keep the economy close to its initial (efficient) equilibrium. As noted earlier, the classic example of such a tax is a head tax: each citizen is simply charged the same figure and sends the money to the government. This theory leads to a number of novel implications. Research shows that, compared to females, male labour supply is relatively inelastic; as a result, taxing males' labour supply is relatively efficient. In addition, Greg Mankiw has a nice—but mainly provocative—paper suggesting that taller people should be taxed more. Idea is, research indicates that taller people have higher incomes (one possible reason, he claims, is that taller people are more self-confident.) Optimal tax theory dictates that taller people should be taxed more; namely, height is a good proxy for income, but *crucially*, you can't change how tall you are (i.e., height is inelastic.) Because people can take action to reduce income, taxing income is inefficient and results in deadweight losses.

4.1 Taxation and Efficiency

According to the standard general equilibrium theory, under certain assumptions, the equilibrium an economy attains is socially efficient. That is, it maximizes the economic surplus—the sum of consumer and producer surplus. Crucially, it is the price mechanism which directs the economy towards this efficient equilibrium. Yet, by introducing taxes, the government distorts the economy away from this equilibrium. If the previous equilibrium was the optimal one, then the one with taxes must be Pareto inferior. Essentially, the tax-adjusted prices direct the economy towards a "wrong" sub-optimal equilibrium. To take a simple example, suppose my willingness to pay for a good is 10. The current price is 8, leading to a consumer surplus of 2. But if the government introduces taxes and raises the price to 11, I don't purchase the good. As such, the taxation has caused a deadweight loss: I lose out, and the government raises no revenue. Instead, I go and buy something which I hadn't bought in the first place (so by revealed preference, I must be worse off.) Although it is socially optimal for this transaction to take place, it doesn't. As another example, suppose the marginal cost of producing two goods is the same: 1. And as is standard, suppose the marginal cost is increasing with the level of production. In addition, suppose consumers value each good equally and utility of each good is subject to diminishing marginal utility. To maximize utility, therefore, the consumer splits consumption over both goods. It is socially efficient for the economy to produce an equal quantity of each good. However, if the government introduces a tax on good 1, it leads consumers to bias their consumption towards the untaxed good. As a result, production will also be biased towards this good. But it is socially efficient for the economy to produce an equal quantity of each: because marginal cost rises disproportionately as more is produced, concentrating production on one good is inefficient. [1]

4.2 Optimal Taxation

To give an example, consider two goods, apples and oranges. There is diminishing marginal utility to both goods. If the price of each good is 1, and the consumer has an income of ten, then he'll purchase 5 of each good. Now suppose the government must raise revenue of 2. To raise this revenue, the government taxes apples but leaves oranges tax-free. But with this large tax on oranges, the consumer would end up purchasing a lot more apples; that is, the consumer is induced to deviate a lot from his optimal plan. Before, he was consuming $(5, 5)$; now he's consuming $(8, 1)$ (say). It'd be better to design taxes so as to entice the consumer to purchase a bundle like his original one. Motivating consumers to do otherwise is suboptimal. In this case, we should design taxes to get the consumer to choose a combo as close as possible to 5, 5—the utility maximizing bundle. One way to do this is to place an equal tax on both goods. In this case, the consumer will choose a bundle, say 4 and 4, that is closer to the optimal one.

Another application of this relates to the intertemporal substitution of labour. Recall from above (assuming $\beta(1 + r_{t+1}) = 1$). More importantly, recall that, since work is painful, the consumer ideally wants to spend an equal amount of each period working. Now suppose there are only two periods in the world, and the wages are equal in both periods; $w_{t+1} = w_t = w$. Then the household's optimal labour

[1] To see this formally, recall from micro that optimal production leads to the relationship $MRT = \frac{P_1}{P_2}$; that is the marginal rate of transformation equals the ratio of prices the firm charges. For the consumer, optimality requires $MRS = \frac{P_1}{P_2}$. Ordinarily, this leads to $MRT = MRS$, which indicates production is socially efficient. Suppose now the government requires that for each unit of good 1 purchased, the consumer must give an extra t to the government. Introducing this tax implies the consumer optimality condition becomes $MRS = \frac{(1+t)P_1}{P_2}$. In turn, this implies $MRT \neq MRS$, so the usual condition for a general competitive equilibrium breaks down (and we can no longer infer Pareto efficiency.)

$$\frac{l_{t+1}}{l_t} = \left(\frac{w}{w}\right)^{\frac{1}{\sigma}} = 1 \qquad \Rightarrow l_{t+1} = l_t.$$

Because work becomes increasingly painful in each period, it is optimal for consumers spread labour supply over two periods. As in the case of consumption, however, the consumer will often find it optimal to deviate from the complete smoothing if relative wages change.

Now suppose the government must raise tax revenue. If we try to collect all the tax in period two, then tax rates would rise then and the wage would fall to $(1-\tau)w$. Returning to the formula above, relative labour supplies become

$$\frac{l_{t+1}}{l_t} = \left(\frac{(1-\tau)w}{w}\right)^{\frac{1}{\sigma}} \qquad \Rightarrow l_t > l_{t+1}$$

In response to the tax, the household would shift labour towards future. But unequal taxes distort the consumer's decision and motivate him to work excessively today. The tax system, therefore, has motivated him to deviate a lot from his initial plan. Yet if the government chose to spread the burden over two periods and set a tax rate $t' < t$ in *each* period, the optimal division of labour would be

$$\frac{l_{t+1}}{l_t} = \left(\frac{(1-\tau')w}{(1-\tau')w}\right)^{\frac{1}{\sigma}} \qquad \Rightarrow l_{t+1} = l_t$$

This is a better outcome, and closer to the worker's optimal decision in the absence of tax. Intuitively, work is painful, so it is better to incentivize the worker to spread labour over two periods. This way, the worker is not killing himself in one period and slacking off in another. Point is, equal tax rates in both periods ensure the worker doesn't kill himself in any period. Formally, we say "tax smoothing" minimizes the distortions associated with taxation. It is easy to show formally that the worker is ultimately better off in this tax regime. An implication of this is the government should run deficits when expenditure is unduly high—and not raise taxes. Instead, they should keep taxes constant over the cycle and use any surpluses to pay off debt.

In some cases taxes are placed on goods like cigarettes *so as to* distort behaviour *away* from consumption of the given good. If a good has a negative externality, then we know that "too much" is produced. For this reason, a tax on the good will distort this initially sub-optimal equilibrium and ideally lead the economy towards the optimal one. Such taxes are called *Pigou taxes.*

4.2.1 Issues in Taxation

Incidence

The *incidence* of a tax refers to those who are ultimately burdened with it. For instance, we showed already that a tax on capital would lower *wages* in the long run, so workers beared some of the burden. In general equilibrium, everything affects everything else, so the incidence of taxation is often unclear.

So far we have dealt with the case of *proportional* taxation; that is, a *constant* tax rate; here, the marginal tax rate—which govern substitution effects— was equal to the average tax rate—which mediate income effects. If, for example, labour was taxed at rate t, then the after-tax wage was $(1-t)w$. Yet in reality taxation tends to be progressive. That is, marginal tax rates rise with the level of income. In this case, a different tax rate would apply to a higher w. Suppose that any income above 40000 is taxed at 40% and anything below at 20%. So if a person (say, John) earning 40000 earns anything else, 40% of that goes to the government. Although John obviously doesn't like this, he might work anyway (he want to buy a home, say). However, if the marginal tax rate rises to 70% say, then John might be deterred from extra work effect; the nicer home isn't worth the effort. As you can see, this kind of progressive taxation is potentially quite destructive. In particular, those on higher income are generally the most skilled—recall that wages reflect marginal product— and these are not the people we want to deter from working. Empirical evidence on this topic indicates that the *reported* incomes of the richest people *are* quite sensitive to higher tax rates. In other words, when their marginal tax rates are reduced, they ultimately pay *more* in tax, and vice versa. However, this could be due to less tax evasion under lower marginal tax rates—and not necessarily less economic activity (but from the standpoint of raising revenue, it doesn't matter.) Figure 1.4 illustrates the Laffer curve. Laffer argued that there was an optimal rate of tax that maximized revenue. In particular, any increases beyond this rate *lowered* revenue. The disarming simplicity of the Laffer curve has made it extremely influential in policy circles—and arguably has been one of the most influential theories in the last 40 years or so. In Ireland, for example, *reducing* the corporate tax rate has almost surely caused tax revenues to *rise*.

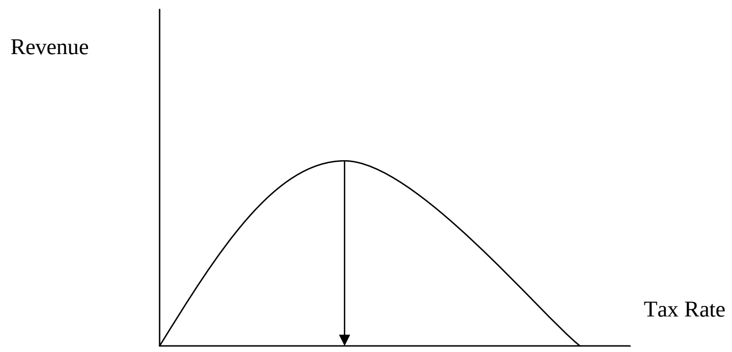

Figure 4.4:

4.2.2 Lump-Sum and Proportional Taxation

To illustrate how lump-sum taxation and tax rates have different economic implications, I present an example. Suppose consumer has utility function

$$u(C_1) + \beta u(C_2) = \log C_1 + \beta \log C_2$$

The consumer faces a real wage of w each period.

The usual intertemporal budget constraint is

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

Now, in this example, the usual labour optimality condition (for period 1) reduces to

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

Lump-Sum Taxation

If we impose a lump sum tax, T, on the consumer in period 1, then the budget constraint is

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

We know from the *permanent income hypothesis* that consumption in each period will be lower. Relative to before, therefore, C_1 and C_2 will be lower, and marginal utilities higher. From the labor optimality condition, $v'(l)$ must be higher (i.e., the pain of work has increased)—but this implies that the consumer is working harder and has increased labor supply. In response to the tax, the consumer reduces consumption of goods and consumption of leisure. In short, the imposition of a lump sum tax makes the consumer feel poor and *raises* labour supply.

Tax Rates

Imposing a proportional (i.e., constant) tax rate t implies the intertemporal budget constraint becomes

$$C_1 + \frac{C_2}{1+r} = (1-t)wl_1 + \frac{(1-t)wl_2}{1+r} = wl_1 + \frac{wl_2}{1+r} - \left(twl_1 + \frac{twl_2}{1+r}\right)$$

Meanwhile, the labour optimality condition becomes

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

The effects now are different. Looking at the budget constraint, notice that the consumers lifetime income is lower than before, which will have similar qualitative effects to the lump-sum tax; lifetime income is lower, so consumption will be lower each period. Yet the labour optimality condition has also changed this time. In particular, $u'(C_1)$ has risen, but the after tax wage, $(1-t)w$, has fallen. Overall, therefore, we don't know what happens to $v'(l)$ and hence labour supply. Namely, in this case, there are both income and substitution effects to the imposition of the tax on labour income. The lump-sum tax, by contrast, did not change the price (i.e., the wage), so did not introduce a relative price distortion.

Finally, recall from real business cycle theory that a *temporary* reduction in tax rates for a period can unambiguously raise labour supply that period; this was the intertemporal substitution of labour. Although we are primarily concerned with the long-run policy here, this is one instance where changing tax rates can be used as a fiscal stimulus.

4.2.3 Types of Taxation

Taxes on Capital

Most economists agree that taxing capital (interest income) is sub-optimal. If you invest W and there is a tax of t on interest income, then at the end of T periods, you will receive $(1+(1-t)r)^T W$. Point is, as T increases, the distortion increases as fast due to compounding. As a result, the decision to save for a long period becomes highly distorted. So with a capital tax, savings becomes increasingly less attractive. Because savings are important for generating funds for investment, this has adverse effects for the economy. Rather than taxing savings and income (which both lead to substantial distortions), many economists advocate a *consumption tax*. Because this is hard to avoid—we all have to consume, and ultimately expect to consume all income—it introduces fewer distortions.[2] One common argument against a consumption tax is its obvious lack progressivity: both rich and poor would ultimately end up paying the same *proportion* of their income in tax.

Means Testing

Any policy whereby some benefit disappears if you earn more has the same implication as an increase in the marginal tax rate. For example, suppose that children's allowance is contingent on how poor you are; and in particular eligibility is contingent on household income below 30000. Therefore, if household income rises above 30000, the household becomes ineligible for child benefit. Clearly, this policy acts just like an increase in the marginal tax rate for those earning 30000. Another prominent example is unemployment insurance: if a person finds a job, they become ineligible. And the higher the unemployment insurance, the greater that tax on entering the workforce. For this reason, many economists advocate keeping unemployment insurance relatively low; unsurprisingly, countries with high levels of unemployment insurance–such as France—have relatively high unemployment rates. Formally, we say that a person will enter the labour force if the wage w exceeds what's called the reservation wage, $\overline{w}$; that is, when $w > \overline{w}$. The reservation wage is the minimum wage required to induce the person to work, and its level is naturally increasing in the level of unemployment insurance. Raising taxes on labour causes w to fall to $(1-t)w$ and thereby makes it less likely that $(1-t)w > \overline{w}$. Similarly, raising unemployment insurance increases $\overline{w}$ and again makes it less likely that $w > \overline{w}$.

[2]One distortion is that it could lead to a substantial underground economy.

Inflation Tax

One way for the government to finance expenditure is to print money. By the quantity theory, this in turn increases the price level, and reduces real income. In this sense, printing money is said to confer an *inflation tax.* You could argue that the inflation tax is efficient: it raises prices and and depresses real wages, raising labour supply (it confers a negative pure income effect.) It does not distort relative prices.

Risk/Entreprenuership

Because taxation reduces the benefits from taking risks (but typically has no effect on the losses), taxation tends to reduce risky activities and entrepreneurship. To see why, consider this example. Suppose that there is a risk project that pays -10 with probability .5 and 12 with probability .5. So the expected gain to this is 1, and assuming the investor is risk-neutral, the project will be undertaken. But now suppose the government taxes gains at 50%. So now the project pays -10 with probability .5 and 6 with probability .5, yielding an expected value of -2. Because taxation reduces the returns from the project, it will not be undertaken. In this sense, taxation deters entrepreneurship, and given the importance of innovation to sustained growth, this is a potentially harmful implication.