The Schumpeterian Model of Growth

Chad Jones and Dietrich Vollrath

Introduction to Economic Growth

Schumpeterian model of growth

A lot like the Romer model, but with key differences in **bold**

- Has all the elements of the Solow model
- Is specific about what A means in that model
- A **represents quality of goods, not number**
- Explains the dynamics of A and g_A
- Explains why g_A is constant along a BGP
- Requires effort to create the new ideas that drive A and g_A
- Explains the choice involved in making that effort
- **Is explicit about the role of competition**

The Solow part

Much of the model starts with familiar items. Production is

$$Y_t = K_t^{\alpha} \left(A_t L_{Yt}\right)^{1-\alpha},$$

where L_{Yt} are workers employed in producing goods and services. L_{Rt} are people engaged in R&D producing ideas, and

$$L_t = L_{Yt} + L_{Rt}. \tag{1}$$

denoting the ratio of R&D workers as

$$s_R = \frac{L_{Rt}}{L_t}. \tag{2}$$

The Solow part

Standard parts:

- Capital accumulates in the same way as in the Solow (depends on K/AL)
- Population growth is the same as in the Solow (exogenous at g_L)

Assuming that g_A ends up constant (which we'll see) then economy ends up at steady state with a constant K/AL ratio as usual.

Schumpeterian ladders

Here the concept of innovation is in improving on an existing idea

- For example, a new version of an iPhone, or a new model of a car
- An innovation is a step “up” the ladder of quality
- The steps are discrete, so innovation is “lumpy”
- R&D will increase the probability of taking a step
- This means actual growth will be lumpy (zero one year, positive the next)
- But the trend growth of GDP per capita will be smooth

Improving an idea

Schumpeterian mechanics for ideas are

$$A_t = (1+\gamma)^{N_t}, \tag{3}$$

The level of productivity at time t depends on

- N_t. The number of steps we've taken up the ladder of quality
- γ, the "step size" of each innovation. Assumed constant.

Taking steps up the ladder

The process for taking steps depends on research effort as in

$$E[dN] = \theta \frac{L_{Rt}^{\lambda}}{A_t^{1-\phi}}, \tag{4}$$

where

- L_{Rt} is the labor doing R&D
- A_t is the existing position on the ladder
- θ scales the process
- λ and ϕ work exactly like in the Romer model

Taking steps up the ladder

A key distinction in this equation

$$E[dN] = \theta \frac{L_{Rt}^{\lambda}}{A_t^{1-\phi}}, \tag{5}$$

is that this is the *expectation* of steps, dN:

- R&D effort increases the probability of taking steps
- And the existing level, A_t, may make it more or less likely you jump

The growth rate of productivity

To get to a growth rate for produtivity, g_A, start with logs:

$$\begin{aligned}\log A_t &= N_t \log(1+\gamma) \\ &\approx N_t \gamma,\end{aligned}$$

and take the time derivative to find

$$g_A = \frac{dA}{A_t} = \gamma dN. \tag{6}$$

The growth rate depends on the step size, and the number of steps.

Expected growth

Because the number of steps is uncertain, so is the growth rate.

$$E[g_A] = \gamma E[dN]. \tag{7}$$

which implies that

$$E[g_A] = \gamma\theta \frac{s_R^{\lambda} L_t^{\lambda}}{A_t^{1-\phi}}.$$

This looks a lot like the Romer model now, but with the distinction of that expectations operator.

What does growth look like?

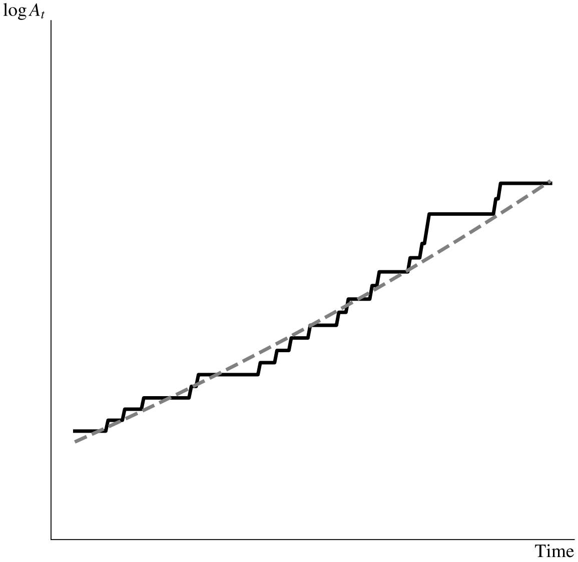

Dynamics of productivity growth

Despite the expectations operator, the long-run growth rate ends up like the Romer model

$$E[g_A]^{ss} = \frac{\lambda}{1-\phi} g_L.$$

In expectation the growth rate still just depends on how fast the population grows.

The role of step size

Note that in this steady state

$$E[g_A]^{ss} = \frac{\lambda}{1-\phi} g_L.$$

The step size of innovation, γ, does not matter at all. Why?

- A big γ means big jumps in productivity, A_t, when innovation occurs. That should make the growth rate higher.
- But that same jump in A_t also makes $E[dN]$ *lower* because it raises the existing level of productivity
- These two effects cancel out in steady state. The rate at which you climb the ladder depends on how much effort you put in, g_L, not on how far apart the steps are, γ.

The role of step size

But note that in steady state the rate of $E[dN]$

$$E[dN]^{ss} = \frac{E[g_A]^{ss}}{\gamma} = \frac{\lambda}{1-\phi}\frac{g_L}{\gamma}.$$

does depend on γ.

- For a given level of effort, if the steps are bigger, the fewer you take.

Comparative dynamics

All of the logic and analysis about changes in s_R, θ, and g_L are the same in the Schumpeterian model as in the Romer

- There is no long-run effect on the growth rate from s_R. It does affect the level of A_t
- The size of g_L changes the long-run growth rate and the level of A_t
- The diagram and analysis of changes is identical to the Romer model

The choice of R&D

Like Romer, all the long-run results hold regardless of the choice of s_R. But s_R determines the level of productivity and GDP per capita. What determines s_R?

- How do firms make the decision to do R&D? Fixed cost versus flow of profits
- What determines the fixed cost?
- What determines the flow of profits?

This requires an explicit description of an imperfect market which allows market power and profits.

The cost of R&D

For a potential firm, the fixed cost of finding a new idea to implement is

$$F_t = w_t \frac{L_{Rt}}{E[dN]}. \tag{8}$$

- w_t is the wage they have to pay to people to do R&D
- $L_{Rt}/E[dN]$ is how many workers it takes *per step*
- Potential firms take this ratio as given, determined in the aggregate
- Hence F is wages/worker times worker/idea = wages/idea

The benefit of R&D

Here is where the Schumpeterian model gets interesting.

- A firm which owns the highest "step" on the ladder will be the one who supplies all the intermediate goods to the final goods provider. We'll show why later.
- A potential firm will do R&D to try and find the next "step", and take over as the supplier.
- This is the "creative destruction" that Schumpeter wrote about, as one firm replaces another
- In Romer, all firms continued to exist forever once created - no destruction
- But the potential firms have to account for the chance that they will be destroyed
- Their eventual replacement means firms value an idea a little less, ceteris paribus

The benefit of R&D

For a potential firm, they can earn profits from an idea but also can be replaced. They care about the present discounted value of those profits, V, and will compare that to F.

$$V_0 = \pi_0 + \frac{\pi_0(1+g_\pi)(1-E[dN])}{(1+r)} + \frac{\pi_0(1+g_\pi)^2(1-E[dN])^2}{(1+r)^2} + \ldots$$

- π_0 is profits today
- r is the rate of return
- g_π is the growth rate of profits
- The $(1-E[dN])$ terms are the chance that the firm will be replaced in the future. Note that this probability goes up over time. Eventually they *will* be replaced.

The benefit of R&D

We can reduce that value of R&D

$$V_0 = \pi_0 + \frac{\pi_0(1+g_\pi)(1-E[dN])}{(1+r)} + \frac{\pi_0(1+g_\pi)^2(1-E[dN])^2}{(1+r)^2} + \ldots$$

to

$$V_0 = \pi_0 \sum_{t=0}^{\infty} \left(\frac{1+g_\pi - E[dN]}{1+r} \right)^t,$$

and then

$$V_0 = \frac{\pi_0}{r - g_\pi + E[dN]}. \tag{9}$$

The innovators decisions

We have lots of potential firms, and so long as $V_0 > F$ they will continue to do R&D. Assume they enter until $V_0 = F$,

$$\frac{\pi_t}{r - g_\pi + E[dN]} = w_t \frac{L_{Rt}}{E[dN]}.$$

Ultimately we want to solve this for s_R, which recall is $L_{Rt} = s_R L_t$. Like before, we are solving for L_{Rt}. But we need:

- Initial profits
- Growth rate of profits
- Wage

Overview

The overall structure of this economy is both harder and easier than Romer.

- At the "top" there are a set of final goods firms (e.g. Target). They stock *one* intermediate goods that consumers purchase (e.g. just Diet Coke)
- Final good firms are competitive (no profits) and there is no innovation here. They stock goods only.
- There are intermediate good firms that compete to supply the single intermediate good.
- Each intermediate firm's product has a value of A that influences the final good firm's productivity (e.g. the product is more valuable to consumers)
- The intermediate firm with the highest value of A will supply the intermediate good

Final good firms

The final good firms produce GDP using

$$Y = L_Y^{1-\alpha} A_j^{1-\alpha} x_j^{\alpha}. \tag{10}$$

- L_Y are production workers (not R&D workers)
- x_j is the amount of the product they stock
- A_j is the *level* of quality of the product they stock
- α captures how important the product is versus labor

Note that this Y is GDP because intermediate good sales to the final good firm are explicitly not accounted for in GDP.

Final good maximization

We assume final good firms maximize profits. They are competitive so profits end up at zero, but they still try. To do this they set marginal product equal marginal cost. For labor:

$$w = (1-\alpha)\frac{Y}{L_Y} \tag{11}$$

and for the intermediate product j

$$p_j = \alpha L_Y^{1-\alpha} A_j^{1-\alpha} x_j^{\alpha-1}. \tag{12}$$

The important feature here is that the final good firm will pay *more* for a higher quality product, A_j, all else equal.

Intermediate good firms

Each intermediate firm produces their good using the function $x_j = K_j$, using only capital, which costs r per unit of capital. Their profits are:

$$\pi_j = p_j x_j - r x_j.$$

ASSUME for the moment that the firm highest on the

technology ladder (N) acts as the sole monopolist. We'll come later to why. They know how p_j responds to their choice of x_j.

That is, they know what the demand curve of final good firms looks like. They set marginal revenue equal to marginal cost

$$p + \frac{\partial p}{\partial x} x = r.$$

Intermediate good firms

Take the MR = MC condition

$$p + \frac{\partial p}{\partial x} x = r.$$

divide by p

$$1 + \frac{\partial p}{\partial x} \frac{x}{p} = \frac{r}{p}.$$

and the ratio on the left is the elasticity of price with respect to quantity, which from final goods firms is equal to $\alpha - 1$. So

$$1 + (\alpha - 1) = \frac{r}{p}$$

and we get the normal markup of

$$p = \frac{1}{\alpha} r. \tag{13}$$

Adding up

Given the "micro" results on final good firms and intermediate firms.

- The best intermediate firm supplies that, $p_N = r/\alpha$
- Because there is one intermediate firm, it uses all capital $x_N = K$

so final good firms produce

$$Y = L_Y^{1-\alpha} A_N^{1-\alpha} x_N^{\alpha}. \tag{14}$$

which solves to

$$Y = K^{\alpha} (A_N L_Y)^{1-\alpha}. \tag{15}$$

Again this looks just like the Solow, with A_N the step on the technology ladder.

Wages and profits

Solve for things we need to know. Start with the wage. Given the first order condition from final good firms:

$$w_t L_{Yt} = (1-\alpha) Y_t,$$

$(1-\alpha)$ of GDP gets spent on workers. The other α must get spent on the intermediate good (final good firms don't have profits). The revenues of the single intermediate firm are thus

$$p_N x_N = \alpha Y_t \tag{16}$$

Profits per firm

What are profits for the intermeidate firm?

$$\begin{aligned}\pi_t &= p_t x_t - r_t x_t \\ &= (p_t - \alpha p_t) x_t \\ &= (1-\alpha) p_t x_t,\end{aligned}$$

and plug in for firm revenues $p_N x_N = \alpha Y_t$

$$\pi_t = (1-\alpha)\alpha Y_t. \tag{17}$$

are profits per intermediate firm. This is the initial profits that go into the valuation of an idea.

Growth rate of profits

How fast do profits grow? Profits are

$$\pi_t = (1-\alpha)\alpha Y_t. \tag{18}$$

so profits grow at the rate (along a BGP) of

$$g_\pi = g_Y$$

but we know along a BGP that $g_Y = g_L + g_A$ so

$$g_\pi = g_L + g_A \tag{19}$$

or profits in the Schumpeterian model grow both with the market and as steps are taken on the ladder.

Solving for s_R

We have all the pieces to go back to this condition for innovation:

$$\frac{\pi_t}{r - g_\pi + E[dN]} = w_t \frac{L_{Rt}}{E[dN]}.$$

so

$$\frac{\alpha(1-\alpha)Y_t}{r - g_A - g_L + E[dN]} = (1-\alpha)\frac{Y_t}{L_{Yt}} \frac{L_{Rt}}{E[dN]}.$$

which solves for

$$\frac{s_R}{1 - s_R} = \frac{\alpha(1-\alpha)}{(1-\alpha)} \frac{E[dN]}{r - g_A - g_L + E[dN]}.$$

Solving for s_R

Given

$$\frac{s_R}{1-s_R} = \frac{\alpha(1-\alpha)}{(1-\alpha)} \frac{E[dN]}{r - g_A - g_L + E[dN]}.$$

and we know that $E[g_A] = \gamma E[dN]$ given the structure of technology, so we can arrange this to

$$\frac{s_R}{1-s_R} = \frac{\alpha(1-\alpha)}{(1-\alpha)} \frac{E[dN]}{r - g_L + E[dN](1-\gamma)}. \tag{20}$$

s_R is higher:

- If r is lower. If the future matters more, it pays to do R&D
- If g_L is higher. If the market will grow quickly, it pays to do R&D
- If $\alpha(1-\alpha)$ (profits as a share of GDP) is higher
- If $(1-\alpha)$ (wages as a share of GDP) is lower

Dual role of innovation

Unlike the Romer there are two distinct effects of innovation, $E[dN]$, here:

$$\frac{s_R}{1-s_R} = \frac{\alpha(1-\alpha)}{(1-\alpha)} \frac{E[dN]}{r - g_L + E[dN](1-\gamma)}. \tag{21}$$

What happens if $E[dN]$ is higher?

- In the denominator, this *lowers* s_R
- This captures that higher $E[dN]$ means higher chance of being replaced
- In the numerator, this *raises* s_R
- This captures that higher $E[dN]$ makes it easier to replace the existing supplier
- On net, the numerator "wins" mathematically
- Economically, the immediate gain of profits from innovation outweighs the distant chance of replacement

Drastic innovation

We assumed that the firm with the best technology was the sole monopoly provider of the intermediate good. Under what conditions does that assumption make sense?

- If the "step size" of innovation, γ, is big enough, this will hold
- The final good firm values the intermediate because of the level of A_j it provides
- The cost of producing the intermediate is the same for firms with different levels of A_j
- The final good firm will pay more for A_N than for A_{N-1}
- The firm with the A_N technology can lower their price until A_{N-1} can't make a profit
- That leaves A_N firm as the only supplier
- This only works if the step size is sufficiently large

Competition

Formally, let the A_N and A_{N-1} firms engage in Bertrand competition. They set prices separately, but understand the behavior of the other, and act strategically. They set their price knowing how the other firm will react. Compare the demand

curves for their products. For a given quantity of x_j, what is the relative price that the final good firm will pay for the better technology

$$\begin{aligned}
\frac{p_N}{p_{N-1}} &= \frac{\alpha L_Y^{1-\alpha} A_N^{1-\alpha} x^{\alpha-1}}{\alpha L_Y^{1-\alpha} A_{N-1}^{1-\alpha} x^{\alpha-1}} \\
&= \left(\frac{A_N}{A_{N-1}}\right)^{1-\alpha} \\
&= (1+\gamma)^{1-\alpha}.
\end{aligned}$$

which depends on how big the step size is.

Strategic interactions

Think through the interaction:

- The lagging firm wants to stay in business. They can try to lower their price and steal the business from the leading firm.
- The lagging firm lowers their price, the leading firm will follow, maintaining a ratio of $(1+\gamma)^{1-\alpha}$ so that the final good firm chooses their product
- The lagging firm lowers again, and the leader follows, and so on
- The lagging firm has to stop lowering when $p_{N-1} = r$, or when the price they charge equal MC and leaves them with zero profits
- We know the leading firm will charge at *most* $p_N = (1+\gamma)^{1-\alpha} r$ because of the lagging firm pushing down the price

Strategy for N

We know N firm will charge at most $p_N = (1+\gamma)^{1-\alpha} r$

- If they charge exactly this then the lagging firm can still sell some units to the final good firm
- if the leading firm lowers the price further, they can take all the business.
- Will the leading firm go lower?
- The leading firm's profit-maximizing price is $p_N = r/\alpha$ as the monopolist
- If $r/\alpha < (1+\gamma)^{1-\alpha} r$ the profits are higher with the lower price
- The leading firm will set $p_N = r/\alpha$ if $(1+\gamma) > (r/\alpha)^{1/(1-\alpha)}$.
- If the step size γ is "drastic", then the leading firm lowers p_N until the lagging firm is out of business.
- The assumption of a single monopolist is an assumption that innovation is "drastic" and that γ is large