

THE ECONOMICS OF INNOVATION

Toronto Rotman PhD Seminar | Spring 2021 | Kevin Bryan

CLASS 1 NOTES

These notes include three sections: first, what does Arrow think is the fundamental problem of research; second, how do we go wrong with common qualitative intuitions like Schumpeter's theory that since bigger firms can better exploit research they will do more of it; and third, how do growth models with endogenous innovation like Romer (1986, 1990), Grossman and Helpman's quality ladder (1991), and Aghion and Howitt (1992) fit in with the classical growth models? I have written these notes such that only calculus is needed to follow the exposition, at the cost of perhaps oversimplifying the original models. These notes are meant to be broadly accessible.

ARROW 1962 AND THE FUNDAMENTAL PROBLEMS OF INNOVATION

Arrow's "Economic Welfare and the Allocation of Resources for Invention" is almost entirely verbal and incredibly prescient. It is not going overboard to say that he lays out the research agenda in the economics of innovation for the next fifty years.

The general equilibrium "welfare theorems" say that if the production function takes commodities and labor and outputs goods in a convex manner, among a few other assumptions, the market equilibrium will be efficient in the Pareto sense. This seems innocuous, but it conceals three assumptions: the production function should be deterministic, the goods should be possible to sell in a market, and the production transformation set should be convex (if I can produce 2 of good x using α of good y, then I need to be able to produce at least 1 of good x using $.5\alpha$ of good y).

Knowledge seems to violate all three of those assumptions. The first unit of a piece of knowledge is costly to create and subsequent pieces of that knowledge are free to replicate; think of a flame being passed from candle to candle costlessly once the first candle is lit. That means I can produce 1 unit of knowledge with α resources, 2 units still only using α , 3 using only α in total, and so on. Note that spending $.5\alpha$ does not mean I can produce at least half as many units of knowledge as I can produce using α , so the production transformation set is not convex, and hence an efficient equilibrium may not exist. Intuitively, efficiency requires that price equal the cost of the marginal unit of knowledge consumed, which is zero. But if the price of knowledge is zero, who will pay the fixed cost to create it?

Even if we could reward firms for the immediate social value created when their knowledge is used, the full value of knowledge goods is not reflected in their market price because there are significant externalities. Principally, knowledge invented today is an input into knowledge invented tomorrow. Even under the strongest patent systems, there is no patent strong enough to cover ideas only *inspired* by the original invention, but the social value of the original invention surely ought include the value of that inspiration. For reasons of individual liberty, firms

may be unable to prevent their trained workers from leaving to new jobs.

The fact that research inputs only produce knowledge with uncertainty is also problematic. In general, the theory of Arrow-Debreu securities says that uncertainty is not a problem for economic efficiency as long as risks can be suitably hedged: with a market allowing proper hedging of risk, “the use of inputs, including human talents, in their most productive mode is not inhibited by unwillingness or inability to bear risks.” But Arrow, very much ahead of his time, noted that the uncertainty of research is partially related to the fact that we can’t observe the effort of the researcher, and hence if there is no risk on the part of the research - that is, if she gets a constant payout no matter whether she succeeds with her invention today or not - then surely the researcher will shirk. The problem of how to balance the efficiency-improving desire to shift risk away from individuals with the incentive problem lies at the heart of much of the 1980s mechanism design theory on organizational incentives which we will cover in depth.

Dealing with these problems is of enormous policy importance. A commentator on Arrow’s paper notes that the military develops planes with production costs in the billions of dollars. How ought they incentivize researching firms? How should risks of the failure of research be shared? Should firms be compensated on the basis of planes eventually bought, or in terms of their research costs, or with prizes, or some other policy altogether? These questions don’t matter in most markets - the invisible hand, as interpreted in Arrow and Debreu, suggests that market prices provide efficient incentives in many cases. Innovation is not one of those cases. To the extent that knowledge is a particularly important good, and we will see in the upcoming discussion of growth that it is, then proper incentives for knowledge creation are all the more important.

HOW MODELS WILL HELP OUR INTUITION

Schumpeter, in his famous *Capitalism, Socialism and Democracy*, argued that monopolies will be better for innovation than competitive firms. His supply side argument was that larger firms can hire more diversified R&D staff, who are more efficient (Axiom 1), and that larger firms are more likely to be able to use the output of quasirandom research since they sell in more product lines (Axiom 2). Lots of people have tested the hypothesis of whether more concentrated markets are more or less innovative. But does Schumpeter’s claim even follow from the axioms?

Let $F(R, N)$ be the dollar value of R&D output per worker, where R are research workers and N are other workers at the firm. $RF(R, N)$, then, is the total dollar of value of R&D output for the firm. Axiom 1 says that average R&D productivity per worker rises with the number of workers ($\frac{dF}{dR} > 0$) and Axiom 2 says that average R&D productivity per worker rises with the size of the firm outside of the research division ($\frac{dF}{dN} > 0$). Schumpeter’s theorem is that total R&D output of the firm ($RF(R, N)$) rises more than proportionally with the firm size ($\frac{d(RF)/RF}{d(R+N)/(R+N)} > 1$). Empirical tests of Schumpeter’s hypothesis often test the related idea of whether the number of R&D workers R rises more than $R + N$ when $R + N$ changes over firms ($\frac{dR/R}{d(R+N)/(R+N)}$).

There are two problems. One, as pointed out by Carlos Rodriguez in the JPE in

1979, Axiom 1 implies increasing returns to hiring R&D workers, which implies that the marginal product of researchers exceeds the average product, which implies that if the market for R&D workers is competitive and hence workers are paid their marginal product, there does not exist a finite wage for researchers. For this reason, let's be charitable to Schumpeter and assume Axiom 3, which is implied by Axioms 1 and 2 instead: that if the number of researchers and other workers grow positively at the same rate, then the average R&D output per worker increases.

Here is where formalism helps. Since firms will hire research workers until their marginal product equals their wage, big profit maximizing firms will only hire more R&D workers than small ones if the marginal product of R&D workers ($\frac{dRF}{dR}$) is larger when the firm is larger. That is, the following derivative must be positive:

$$\frac{d}{d(R+N)} \frac{dRF}{dR} = \frac{d}{d(R+N)} \left(F + \frac{dF}{dR} R \right) = \frac{dF}{d(R+N)} + R \frac{d^2F}{dRd(R+N)}$$

That is, the Schumpeterian conclusion depends on the sign of the *second* derivative of the production function, not the first: economic behavior is driven by marginal properties, not average properties. That is, the Schumpeter axioms do not imply the conclusion being tested in so many empirical papers!

You may wonder, who cares? Don't we just care about the empirical question of whether monopolies do more research or not? Well, Schumpeter's qualitative theory suggests a *mechanism*, and further suggests that if you find out that monopolies do more research, it is because they are monopolies, and not for some other reason. Hence you would be justified in offering policy advice like "allow small firms to merge so that they will do more research". Since the theoretical logic is proven to be incorrect when we write things down formally, all we can say from an empirical result of that type is that monopolies do more research than other firms but we have no idea whether it would be good policy to allow smaller firms to merge.

(Beyond the theory, we will see later that it is likely not true that monopolies do the most research. But hopefully this is enough to convince you that looking at ideas in innovation formally can be of benefit even if you are empirically minded.)

CLASSIC MODELS OF GROWTH

What causes economies to become rich? If you are in this class, you likely buy the idea that innovation - new knowledge, widely diffused - is essential. If this were not true, the topics covered in this course would be of much less interest.

This idea, that knowledge is produced by profit-maximizing firms, and that a large knowledge base means a rich country, have proven incredibly difficult to generate in a model. Two other stylized facts, that as population grows the growth rate increases (the past two decades have been the fastest growth in global GDP in human history) and that as the number of researchers grows with population we do not see ever-increasing growth rates (U.S. productivity growth is fairly constant despite a massive increase in the number of researchers over time) would also be nice features for a workhorse model of growth to generate. Let's first see how the classic Harrod-Domar and Solow models think about growth, then see how Paul Romer and his contemporaries' "endogenized" growth.

First, Harrod-Domar. Let the output of the economy $Y = f(K)$ where K is some summed up value of physical capital like machines and steel. Let the economy have constant returns to scale ($\frac{df}{dK} = c$ for some fixed value c). Let output with no capital be zero ($f(0) = 0$). And let the change in capital be the proportion of output saved minus the capital that depreciates ($\dot{K} = sY - \delta K$).

The first three assumptions imply that $y = cK$, therefore the change in output over time $\dot{Y} = \frac{d(cK)}{dt} = c\dot{K}$. Combining those equations, the change in capital over time $\dot{K} = sy - \delta \frac{Y}{c}$, therefore $c\dot{K} = Y(sc - \delta)$, therefore $\dot{Y} = Y(sc - \delta)$, and the growth rate in GDP $\frac{\dot{Y}}{Y} = sc - \delta$. That is, if you want to grow, jack up the savings rate. This model has no role for innovation since the state of technology isn't even the model, nor is there any role for human capital (which, we shall see, is different from technology). This type of growth is sometimes called growth by "capital deepening". The Maoist Chinese policy of "backyard steel production" absolutely fits an idea of a world where one grows by sacrificing consumption for savings, limitlessly.

Solow's famous model brings technology into the picture. Let output be

$$Y = f(AL, K) = K^\alpha (AL)^{1-\alpha}$$

where A is the "state of technology", L is the amount of labor, and K is capital. The production function is constant returns to scale: doubling all factors doubles output. Constant returns to scale at the level of the economy can be justified for now under the "Kaldor fact" that labor and capital shares have been relatively constant for a very long time (or via an argument about replicability - do you see why?). Let labor grow exogenously at rate n ($\frac{\dot{L}}{L} = n$) and let technology grow exogenously at rate g ($\frac{\dot{A}}{A} = g$). As in Harrod-Domar, let capital change depending on savings and depreciation ($\dot{K} = sY - \delta K$), with the savings rate being an exogenous fixed percentage of total output. More complex "neoclassical" models of growth based on Solow would model the households maximization problem directly, but the basic insight of the model is unchanged by that modification.

To solve this model, let output per effective unit of labor

$$y \equiv \frac{Y}{AL} = \frac{K^\alpha (AL)^{1-\alpha}}{AL} = \frac{K^\alpha}{(AL)^\alpha} = k^\alpha$$

The change in capital per effective unit of labor $k = \frac{K}{AL}$ is

$$\dot{k} = \frac{\dot{K}}{AL} - \frac{K}{(AL)^2}(\dot{A}L - \dot{L}A) = \frac{\dot{K}}{AL} - \frac{K}{(AL)}\left(\frac{\dot{A}}{A} + \frac{\dot{L}}{L}\right) = sy - k(n + g + \delta)$$

where the first equality is merely the quotient rule and the final equality comes from the transition equation for capital given above. By this transition equation, the amount of capital per effective unit of labor grows if $sy > \delta k$, shrinks if $sy < \delta k$, and is at a "steady state", neither shrinking nor growing, when $sy = \delta k$. When k is constant ($\dot{k} = 0$), we have that steady state capital per effective unit of labor $k = \left(\frac{s}{n+g+\delta}\right)^{\frac{1}{1-\alpha}}$. Since y is constant when k is, in the "steady state", output per worker $\frac{Y}{L} = Ay$ grows at the exogenous growth rate of technology g .

That is, in the short run society can become richer via capital deepening, but in the long run once steady state capital per effective unit of labor has been reached, only technological growth can push the economy forward. The famous “Solow residual” regression based on this model suggests that for countries at the technological frontier, the overwhelming majority of growth is driven by changes in a residual A , and not by capital or labor increases. This is often interpreted as “productivity” or “technological advances”, but of course controlling for things like human capital is important if that interpretation is desired.

Note the following interesting implication of Solow. If the growth rate of technology g is higher, for every level of technology steady state capital is lower (capital has decreasing returns to scale), and hence income per worker is lower since y is proportional to k . That is, no sense saving so much today if we will be more productive with what we saved tomorrow.

Solow’s model at least allows us to perform exercises like “how much growth is due to changes in capital versus growing populations or growing technology?” Human capital can fairly easily be added to the model as well. However, we are at a bit of a loss trying to endogenize technology. That is, we cannot explain why technology grows or how policy might affect the growth rate. This is unsatisfying for students of innovation: surely the nature of how A changes is fundamental!

Question: What is the only possible growth rate in income per capita in the long run if there is no exogenous increase in technology g ?

ENDOGENOUS MODELS OF GROWTH

If we want growth to increase *endogenously*, then firms must have the incentive to make costly investments in research *and continue to do so* without the economy reaching a fixed point of the level of technology. A great model of endogenous growth would also match a number of stylized facts: potential growth rates are higher in technologically lagging countries than frontier countries, R&D is done by incumbents and by entrants, the link between the amount of competition and the amount of R&D is not monotonic, and so on.

The models essentially come in two classes, one dating to Alfred Marshall and one to Edward Chamberlin. The Marshallian idea is that research has spillovers, so constant returns to scale technology for the firm production functions have increasing returns to scale at the aggregate level. The Chamberlin idea is that there exist markets without perfect competition, attempts to earn rents in those markets provide the incentive to do R&D, and that R&D lowers the cost of future R&D allowing growth to proceed. It is easy in a partial equilibrium IO model of perfect competition to generate incentive for investment: we will see patent race models and the like. What turned out to be hard was writing general equilibrium models where these incentives were neither “too strong” (with growth exploding over time) or “too weak” (with steady state growth of zero).

Some bad news: these models aren’t trivial mathematically. To quote Paul Romer back in 1994, “Dupuit wrote about new goods more than 150 years ago. Economists are inundated with new goods in their daily lives. It is therefore somewhat puzzling that the potential for new goods still plays such a small role in aggregate

economic analysis. In its discussion of the deeper implications of newness, this paper outlines two different forces that may have tended to keep newness in the background. The most obvious restraining force is the technical difficulty of constructing economy-wide mathematical models with fixed costs. The importance of mathematical difficulty has been noted before. [See, for example, the introduction in Krugman (1990) or the initial sections of Romer (1991).] New goods, fixed costs, and market power are relatively easy to capture in a partial equilibrium model, but much harder to incorporate in analysis conducted at the level of the economy as a whole.” The technical difficulty is that ideas, being nonrival, imply increasing returns to scale somewhere in the economy. This is Arrow’s nonconvexity problem. If all factors with constant returns to scale are paid their marginal product, as they must be under perfect competition, there is simply no income left to pay for the factor with increasing returns, and therefore no reason to pay for ideas.

The first successful attempt at incorporating increasing returns into a general equilibrium model comes from Paul Romer’s thesis work (1986); Chad Jones’ 2005 chapter in the Handbook of Growth gives details on earlier failed attempts. Take the Solow model and assume that there exist a measure λ of firms, that capital K is the sum of capital used by all firms, that labor L is likewise, and that the state of technology A is proportional to the amount of capital in society ($A = \gamma K$); that is, there is something like learning by doing going on. Further assume that there is no growth rate in population ($n = 0$). These assumptions mean that there is an externality when firms use capital: each firm’s capital use increases overall knowledge in society, an effect which is not captured by the profits of individual firms. This model, in the spirit of Marshall, generates a growth path with constant increases in capital, production, and technological growth as long as suitable restrictions on the production function and depreciation rates are satisfied (I’ll save you the differential equations...). There is a problem, however: the Romer model has a “scale effect” where population growth generates more capital use generates higher growth rates, counterfactually implying that growth in ever-increasing in population, eventually reaching a growth rate of infinity (see Jones AER 1999)! Further, the 1986 model does not explicitly model knowledge production, hence is unable to serve as a vehicle for studying policies which change the incentives to produce knowledge.

That said, the 1986 paper did inspire a larger endogenous growth literature which explicitly models the incentives for private actors to produce knowledge. Three classics are Romer’s (1990) model of intermediate product variety with market power for the intermediate good generating quasirents, Grossman and Helpman’s model of “quality ladders” where firms have temporary market power via improvements in the quality of a product, and Aghion and Howitt’s (1992) “neo-Schumpeterian” model where invention tomorrow supersedes inventions today.

A simple version of Romer’s “lab-equipment model” is as follows. Let $N(t)$ represent the number of “varieties” of intermediate goods available at any time, and let aggregate production $Y(t)$ depend on a function of those varieties plus labor, with production being constant returns to scale in the intermediate goods and la-

bor. Assume no population growth as in Romer's 1986 paper. In particular, let

$$Y(t) = \frac{1}{1-\beta} X(t)^{1-\beta} L^\beta$$

where, letting $\epsilon_\beta \equiv \frac{1}{\beta}$ be the elasticity of substitution between various intermediate inputs, the intermediate good aggregate $X(t)$ is such that

$$X(t) = \left[\int_0^{N(t)} x(v, t)^{\frac{\epsilon_\beta - 1}{\epsilon_\beta}} dv \right]^{\frac{\epsilon_\beta}{\epsilon_\beta - 1}} = \left[\int_0^{N(t)} x(v, t)^{1-\beta} dv \right]^{\frac{1}{1-\beta}}$$

This functional form means that increases in varieties permit more efficient production of the intermediate aggregate.

The economy can produce consumption goods C , intermediate goods x , or research Z , under the resource constraint $C(t) + X(t) + Z(t) = Y(t)$ and the assumption that all intermediate goods cost ϕ (in terms of final goods) to produce. The economy starts with $N(0) > 0$ varieties, and produces varieties according to $N(t) = \eta Z(t)$; that is, if firms in aggregate spend $Z(t)$ of the final good on research, they will produce $\eta Z(t)$ new varieties. Production of new varieties happens with free entry, production of the final good Y is competitive, all factor markets are competitive, and the inventor of a given variety x has a perpetual patent on that variety allowing the inventor to earn quasirents. The critical features are twofold: invention is done purposefully by firms, and inventors earn quasirents from their inventions.

Since the final good $Y(t)$ is produced competitively, intermediate goods are purchased given intermediate rental prices $p_x(v, t)$ to maximize total production minus the summed rental price plus the wage bill:

$$\max_{x(v,t), L} \frac{1}{1-\beta} \left[\int_0^{N(t)} x(v, t)^{1-\beta} dv \right]^{\frac{1}{1-\beta}} L^\beta - \int_0^{N(t)} p(v, t) x(v, t) dv - w(t)L$$

The first order condition gives that $x(v, t) = p_x^{-\frac{1}{\beta}} L$. By the usual markup rule, the price of each variety is a markup over the cost of production ϕ such that $p^x(v, t) = \frac{\phi}{1-\beta}$. Normalizing $\phi = 1 - \beta$, the price of every intermediate good is 1. Therefore, plugging price into the demand for intermediate goods, we have $x(v, t) = L, \forall v, t$ and profit for each variety at each time t is $\pi(v, t) = 1 \times L - \phi \times L = \beta L$. Finally, substituting the demand for each intermediate good into the final good production function gives $Y(t) = \frac{1}{1-\beta} N(t)L$. Note that there are *constant returns to scale* from the perspective of the final good producer (who takes $N(t)$ as given), but increasing returns to scale for the economy at large: doubling the number of workers and the quantity of each intermediate good doubles final good production, but doubling those factors plus doubling the number of varieties more than doubles final good production.

Taking the first order condition of the final good producer again, it can similarly be shown that the wage $w(t) = \frac{\beta}{1-\beta} N(t)$. Finally, free entry into research means that $\eta V(v, t) = 1$, where V is the lifetime discounted stream of profits from inventing η new varieties at a resource cost of 1 unit of the final good.

Can this model generate growth? If the interest rate is constant, monopolists in each variety make a profit equal to their discounted flow profits, or $V(t) = V = \sum_{n=0}^{\infty} \beta L(1-r)^n = \frac{\beta L}{r}$. Combining this with the free entry of researchers condition, we have that $r = \eta\beta L$. Growth in consumption, and hence growth in overall output, is proportional to the interest rate along a balanced growth path for many representative household utility functions: this is the idea of “Euler condition”, where households are only willing to invest if the interest rate implies that, adjusted for consumers’ time preference, consumption today and consumption tomorrow provide the same marginal utility. Therefore, consumption growth and hence growth in production overall, is proportional to $r = \eta\beta L$. Note again that there are scale effects: if population were growing, growth would explode. Attempts to extend Romer’s model to fix this problem generally involve some sort of “fishing out”, where ideas are harder to produce as the existing number of inventions increases. Note also that the outcome is not Pareto efficient: intermediate goods are marked up, hence used less extensively than they would be otherwise, hence total production is lower than it otherwise would be if inventions were produced by a social planner.

Question: could we regain efficiency if competitors were able to produce the intermediate goods, perhaps at a higher cost than the inventor? This would lower the monopoly markup, but what would be the overall effect in equilibrium?

This model is canonical because it tractably allows analysis of how purposeful market invention can generate growth that is never infinite nor zero in the long run. It is less satisfying because the form of competition among intermediate good producers (aka, inventors) is trivial: they are all monopolists, hence there is no strategic interaction, hence there is no way to investigate issues like the impact of competition on growth.

1 Neo-Schumpeterian Models

How do Schumpeterian growth models like Aghion and Howitt (1992) or Grossman and Helpman (1991) differ (despite the dates, to my knowledge the Aghion-Howitt working paper was written first)? In first generation Romer models, all varieties are always used, and inventors from the past continue to earn the same flow profit. In neo-Schumpeterian models, there is creative destruction: invention today makes some old varieties worthless, meaning that the state of technology, and assumptions about how R&D generates catchup, can be critical. Further, strategic interaction will surely play a role. First, whether a firm invests depends on whether it is destroying its own varieties (a la Arrow’s replacement effect) or destroying rival varieties, and second, the length of time a firm has market power for their inventions depends on the level of inventive effort other firms exert trying to invent an improvement (the “escape competition” effect).

Let’s look at a simple model of this sort, due to Acemoglu, which is easy to compare with the Romer model. As in Romer, the resource constraint means that output is made up of consumption, production of intermediate goods, and research, with intermediate goods produced at cost ϕ . However, instead of a fixed number $N(t)$ of varieties, there is a measure 1 of varieties denoted by $v \in [0, 1]$, where the

“quality” of machine v is represented by a “quality ladder” such that

$$q(v, t) = \lambda^{n(v,t)} q(v, 0), \lambda > 1$$

where n denotes the number of “steps” made on machine v by time t . That is, innovations are improvements in the quality of each machine, with machines leading to proportional increases in quality (though see Kortum (1997) for an argument against this assumption). The final good production function is as in Romer

$$Y(t) = \frac{1}{1-\beta} X(t)^{1-\beta} L^\beta$$

with the intermediate good production function now being

$$X(t) = \left[\int_0^1 q(v, t) x(v, t)^{\frac{\epsilon_\beta - 1}{\epsilon_\beta}} dv \right]^{\frac{\epsilon_\beta}{\epsilon_\beta - 1}} = \left[\int_0^1 q(v, t) x(v, t)^{1-\beta} dv \right]^{\frac{1}{1-\beta}}$$

The implicit assumption here is that only the “leading quality” invention is ever used, no matter its price compared to older vintages, though of course many models relax this feature (see Goettler and Gordon (2011) for a good empirical model of this form).

Steps up the quality ladder are produced in a similar manner to Romer varieties: a firm that spends a flow of $Z(v, t)$ units of the final good on researching line v will generate improvements at flow rate $\frac{\eta Z(v, t)}{q(v, t)}$. This means that improvements are harder for more advanced varieties. Again, there is free entry and the cost of invention is identical for incumbent monopolists and entrants. Note that due to the replacement effect, the entrant always has more incentive to innovate than the incumbent monopolist, and hence all innovation will be done by entrants in a given variety.

Solving for the balanced growth path equilibrium, as in Romer the first order condition for the final good

$$\frac{\partial}{\partial x(v, t)} \left[\frac{1}{1-\beta} \left[\int_0^{N(t)} q(v, t) x(v, t)^{1-\beta} dv \right]^{\frac{1}{1-\beta}} L^\beta - \int_0^{N(t)} p(v, t) x(v, t) dv - w(t)L \right]$$

implies that demand for variety v is $x(v, t, q) = \left(\frac{q(v, t)}{p(v, t)} \right)^{\frac{1}{\beta}} L$. Again, by isoelasticity, the price of the intermediate good is a markup over costs divided by the elasticity of demand. If cost parameter ϕ is normalized to $1 - \beta$, as before, price is optimally $q(v, t)$, as in Romer $x(v, t|q) = L$, and hence $Y(t) = \frac{1}{1-\beta} Q(t)L$ where $Q(t) = \int_0^1 q(v, t)$. Flow profits for inventors are thus $\pi(v, t|q) = \beta L q(v, t)$, which is identical to Romer aside from the quality ladder multiplier.

Translating flow profits to its infinite discounted value, however, is slightly more complex than in Romer, since rivals may replace your invention. The time at which replacement occurs is determined in equilibrium. If you spend 1 unit of the final good in research for a machine at current quality $\frac{q}{\lambda}$ (so that the next invention is a machine of quality q), you will succeed in inventing that next step at flow rate $\frac{\lambda}{q} \eta$. Therefore, by free entry of researchers, $1 = \frac{\lambda}{q} \eta V \Rightarrow V = \frac{q}{\lambda \eta}$. In the balanced growth path, investment on each line is constant and identical (call this z).

Writing down a Bellman equation, where $V(v, t, q)$ is the lifetime discounted value of an invention of quality q , we have that

$$rV - \dot{V} = \pi - zv$$

where $\pi(q, v, t) = p(q, v, t)x(q, v, t) - \phi qx(q, v, t)$ represents the flow profit of having a leading edge invention (don't worry if the origin of the so-called HJB equation is non-obvious; if you don't know this branch of mathematics, just take that formula as given). If the value function of invention is unchanging ($\dot{V} = 0$, as must be the case on a balanced growth path), we can rearrange the Bellman to get

$$V = \frac{\beta qL}{r + z}$$

and combining with the free entry of researcher condition, we have that

$$r + z = \lambda q\beta L$$

As we noted in Romer, nearly any sensible model of consumer preferences will generate an Euler condition which implies that higher interest rates mean higher increases in consumption each period. Under a certain form of these preferences which is easy to work with, the interest rate r is simply the growth rate in production g minus the pure rate of time preference (call this ρ , hence $r = g - \rho$). All that remains is to know how z , total inventive effort, is determined. Intuitively, in a very short period of time Δ , if inventive effort along all research lines is z , then $z\Delta$ inventions will improve one step; if the period of time is sufficiently short, we can ignore the possibility of improving multiple steps along a single variety (formally, this can be shown using what is called Big-O versus little-o orders). Therefore, in Δ , quality will jump λ for a measure $z\Delta(t)$ of the varieties, and not change for the rest. Taking the limit of Δ to zero, then, average quality Q is changing such that

$$\dot{Q} = (\lambda)z - 1 \times z = (\lambda - 1)z$$

Since we have shown that $Y(t) = \frac{1}{1-\beta}Q(t)L$, hence the growth rate in production is linear in the growth rate of average quality, we have that the growth rate in production $g \equiv \frac{\dot{Y}}{Y} = (\lambda - 1)z$. Combining our equations for the interest rate and the inventive effort on each variety with the condition $r + z = \lambda q\beta L$, we have that the growth rate

$$g = (\lambda\eta\beta L - \rho)(\lambda - 1)$$

Note the difference with the Romer (1990) model, where growth was simply proportional to $\eta\beta L$. Here, growth depends on the size of the quality improvement as well. These models are particularly useful for examining how policies affect V - the value of an invention to incumbents - compared to g , the overall growth rate. They can also be expanded, by changing the free entry condition, to permit investigation of various forms of competition in the research sector.