

Quality Ladders, Competition and Endogenous Growth¹

Michele Boldrin and David K. Levine

Department of Economics, Washington University in St. Louis

First Draft: July 10, 2006

This Draft: February 6, 2009

Abstract: We examine a theory of competitive innovation in which new ideas are introduced only when diminishing returns to the use of existing ideas sets in. After an idea is introduced, the knowledge capital associated with that idea expands, and its value falls. Once the value falls far enough, it becomes profitable to introduce a new idea. The resulting theory is consistent with fixed costs of innovation and it accounts for the same facts as the existing theory of monopolistic innovation. However, there is evidence that innovation frequently takes place in the absence of monopoly power and that it is driven by diminishing returns on existing ideas – two facts that the existing theory cannot account for.

¹ We are indebted to years of discussion with Chad Jones, Larry Jones, Robert E. Lucas, Joseph Ostroy, Paul Romer, and James A. Schmitz about these issues. Pete Klenow, Rafael Repullo and Nancy Stokey made many useful comments. We are grateful also to financial support from National Science Foundation Grant No. SES-03-14713 and the Spanish MES under Grant No. SEJ 2005-08783-C04-01

1. Introduction

The standard view of innovation is that progress takes place along a quality ladder – driven by the short-term monopoly power innovators acquire at each step, as short-term monopoly power is essential to compensate for the fixed cost of innovating. For nearly sixty years, since it was first advanced in Schumpeter [1942], this theory has been the primary tool accounting for the dependence of technological progress on economic fundamentals such as patience and cost. Important recent examples of this line of research are the models of Romer [1990], Grossman and Helpman [1991], and Aghion and Howitt [1992].

Our own examination of innovation² suggests that a different story is worthy of consideration. No doubt, each innovation opens the door to moving on to a new rung of the quality ladder. However, when innovation is costly, it is socially and privately best to exploit the adoption and imitation opportunities available on the current rung before moving on to the next. Only when these are exhausted does it become socially and privately optimal to introduce a new innovation. In such a process fixed costs and monopoly power play no essential role.

There are abundant examples to illustrate the different predictions of the two theories, the most obvious being that in the standard theory without monopoly power innovation is impossible, while – as occurs in practice – it is possible in ours. More relevant to the quality ladder framework, the standard view predicts that, for example, after the radio was introduced inventors would move at once to inventing television. Our theory predicts that they would continue to spend their time and energy improving and expanding the production of radios and, especially, their adoption in new profitable venues. Only after radios become widespread, and the gains to their further improvement and expansion driven down, would potential innovators move on to work on television. Our prediction is an accurate description of the facts of the invention of radio and television – and the history of R&D shows it is the rule and not the exception. After the successful introduction of a new product, potential innovators do not turn immediately to inventing something else that may replace it, but rather to imitating and possibly improving the recent invention. After inventing the light bulb, Thomas Edison turned

² See, for example, our recent book: Boldrin and Levine [2008b].

primarily to investing in and promoting electrical power and selling light bulbs – not to inventing the fluorescent light or the LED. His would-be competitors did likewise. As every movie producer can tell you, the time to release your great new blockbuster adventure movie is not two days after your rival has done the same.

The theory of competitive innovation is not new. Aside from the classical theorists – Schumpeter [1912], von Neumann [1937], Plant [1934] and Stigler [1956], among others – recent important contributions include the unpublished works of Funk [1996] and Quah [2002], the paper by Hellwig and Irmen [2001], and our own work, Boldrin and Levine [1997, 2002 and 2008a]. Examples of papers in which competitive innovation is applied to substantive questions other than endogenous growth are, Andolfatto and McDonald [1998], Boldrin and Levine [2001], Funk [2008], Hopenhayn [2006], Prescott and McGrattan [2008]. During the last twenty years, many other authors have contributed to the competitive theory of innovation and growth by investigating various of its aspects, such as the connections between constant returns to scale and economic growth (Jones and Manuelli [1990], Rebelo [1991]), the link between the degree of appropriation and competition (Makowski and Ostroy [1995, 2001]), and the role of indivisibilities in economic development (Acemoglu and Zilibotti [1997]).

The goal of our earlier work was largely to demonstrate the possibility of competitive innovation by clarifying the role of limited capacity in determining competitive rents. Our goal in this paper is to advance the theory to empirical relevance by embodying it in a benchmark growth model that – by capturing the idea that advances are not desirable until existing opportunities have been exploited – can help account for the facts of innovation.

2. The Grossman-Helpman Model

There are a variety of models of quality ladders with fixed costs, increasing returns, and external effects, most notably those of Romer [1990], Grossman and Helpman [1991], and Aghion and Howitt [1992]. We adopt the model of Grossman and Helpman [1991] as a particularly clean example that leads to a simple closed form solution and includes a straightforward welfare analysis. Here we summarize their results, employing their notation throughout.

Goods come in different qualities.³ Denote by d_j the consumption (demand) for goods of quality j , let ρ be the subjective interest rate, and let $\lambda > 1$ be a constant measuring the increase in quality as we move one step up the quality ladder. We let

$$c_t = \sum_j \lambda^j d_{jt}$$

denote quality adjusted aggregate consumption. Utility of the representative consumer is

$$U = \int_0^\infty e^{-\rho t} \log[c_t] dt.$$

One unit of output of each quality requires just one unit of labor to obtain. The first firm to reach step j on the quality ladder is awarded a legal monopoly over that technology. This monopoly lasts only until there is a new innovation and technology $j + 1$ is introduced, at which time all firms have access to technology j . This is the same device used by Romer and by Aghion and Howitt, and has an obvious convenience for solving the model. Taking labor to be the numeraire, the implication is that the price of output of technology $j + 1$ relative to that of technology j is given by the limit-pricing formula $p = \lambda$.

The intensity of R&D for a firm is denoted by \tilde{t} , and the probability of successfully achieving the next step during a period of length dt is $\tilde{t}dt$ at a cost of $\tilde{t}a_I dt$.

Let E denote the steady state flow of consumer spending. Since the wage rate is numeraire and price is λ the monopolist gets a margin of $\lambda - 1$ on each unit sold. His share of E is therefore his margin divided by the price $(\lambda - 1)/\lambda = 1 - 1/\lambda$. Since the cost of getting the monopoly is a_I , the rate of return is $(1 - 1/\lambda)E/a_I$. However, there is a chance ι of losing the monopoly, reducing the rate of return by this amount. Equating this net rate of return to the subjective interest rate gives the Grossman and Helpman equation determining research intensity

$$\frac{(1 - 1/\lambda)E}{a_I} - \iota = \rho.$$

There is a single unit of labor,⁴ the demand for which comes from the $a_I \iota$ units used in R&D and the E/λ units used to produce output. Consequently, the labor resource constraint is

³ In the original Grossman-Helpman paper there were a continuum of identical sectors indexed by ω . Since this plays little role in the analysis, and for notational simplicity, we omit it here.

$$a_I \iota + E / \lambda = 1.$$

Notice that the same labor is used for R&D as is used to produce output. This captures the sensible – and in this model crucial – idea that there is an increasing marginal cost of R&D. That the increasing cost is due to resources being sucked out of the output sector is analytically convenient. It implies that the cost of R&D, measured in units of output, is proportional to the current rung on the quality ladder, making possible steady state analysis.

These two equations can be solved for the steady state research intensity

$$\iota = \frac{(1 - 1/\lambda)}{a_I} - \frac{\rho}{\lambda}.$$

By contrast the social optimum research intensity is derived by calculating steady state utility to be $[\log E - \log \lambda + (\iota / \rho) \log \lambda] / \rho$. Since the optimal plan in a steady state maximizes the steady state utility subject to the resource constraint, simple algebra gives the optimum

$$\iota^* = \frac{1}{a_I} - \frac{\rho}{\log \lambda}.$$

3. Climbing the Ladder under Competition

Innovation is driven by the rush of firms to get ahead of one another, suggesting that competition is the driving force. In the account given by the standard theory, though, fixed cost and monopoly power play key roles. In particular, if the fixed cost of research intensity, a_I , goes to zero, both the equilibrium and the optimal research intensities go to infinity. Conversely, if monopoly power goes to zero so does research intensity, and economic growth along with it.

Diminishing returns also play a role in the model. As R&D increases, it becomes more costly relative to output, while its benefits do not increase correspondingly. These diminishing returns determine both the equilibrium and optimal research intensities. In the equilibrium it is monopoly power that, by overcoming the obstacle imposed by fixed cost, makes research intensity positive allowing both innovation and growth.


⁴ We have simplified Grossman and Helpman's notation by normalizing the stock of labor to one.

No doubt diminishing returns in the R&D sector are part of the reason that we do not move up the quality ladder more quickly. But is it the only reason, or even the most important reason? Consider the example from the introduction of radio and television. Suppose that the technology of radio has just been invented. Do we expect that, if there were no monopoly, the market would be instantly flooded with radios, and an immediate effort put forward to invent TVs? Or is it more reasonable to suppose that, initially, only a few radios could be produced and, due to their scarcity, sold at a fairly high price? Because the price of radios would initially be high, resources will be allocated to expanding their production and not to inventing television. In our story, as radio production ramps up, the price of radios diminishes. It eventually becomes so low that it makes sense to invest resources to inventing and producing TVs. Standard models of innovation make the convenient assumption that output on a given rung jumps immediately to the long-run level. By doing so they cannot account for the fact that diminishing returns on a rung play a key role in actual innovation.

The story of endogenous innovation on a quality ladder we want to tell is one where the incentives for innovating come about because of diminishing returns to making use of previous inventions. Because inventing a new good costs more than reproducing an existing one, it becomes profitable to invent the new good only when the available quantity of the old good is large enough to make its price low, relative to that of the new one.

Before presenting a model of endogenous innovation due to diminishing returns, it is useful to look at a typical example of how real quality ladders work, borrowed from Irwin and Klenow [1994]. The good in question is the DRAM memory chip. The different qualities correspond to the capacity of a single chip. The figure below, showing shipments of different quality chips, is reproduced from that paper. The key fact is that production of a particular quality does not jump up instantaneously but ramps up gradually, and that a new quality is introduced when the stock of the old one is fairly large. Further, the old vintage is phased out gradually as the new one is introduced. Their price data shows that the price of each vintage of chip falls roughly exponentially over the product cycle – meaning that the incentive to introduce the next generation chip keeps increasing. This vividly portrays our story. Evidence suggests this is the usual pattern in

most industries.⁵ The question this evidence poses, and the intuition upon which our model is built is: why introduce a new product if the old one is still doing so well?



4. Innovation with Knowledge Capital

We adopt the same demand structure as Grossman and Helpman, that is quality adjusted consumption is

$$c_t = \sum_j \lambda^j d_{jt}$$

and the preferences for the representative consumer are

$$U = \int_0^\infty e^{-\rho t} \log[c_t] dt.$$

Unlike Grossman and Helpman, we assume that consumption is produced both from labor and the existing stock of specialized productive capacity. For simplicity we identify “productive capacity” with knowledge and assume that different rungs on the quality ladder correspond to different qualities of capital and knowledge used to produce

⁵ See, for example, Hannan and McDowell [1987], Manuelli and Seshadri [2003], Rose and Joskow [1990], Sarkar [1998]

that particular consumption good. We denote by k_j the combined stock of capital and embedded knowledge that goes into producing quality j output. By explicitly modeling the stock of knowledge, we can distinguish between investment on a given rung – spreading and adopting knowledge of a given type through teaching, learning, imitation, and copying – and investment that moves between rungs – innovation or the creation of new knowledge. We refer to k_j as *quality j knowledge capital* or, simply, *knowledge*. In practice this can have many forms – it can be in the form of human knowledge or human capital, but it can also be embodied in physical forms such as books, or factories and machines of a certain design. In our theory it plays the crucial role of keeping track of progress on a particular rung of the quality ladder.

Knowledge has two uses: it can be used either to generate more knowledge or to produce consumption. More knowledge is produced either increasing the stock of the same quality of knowledge capital or creating a higher quality. If quality j knowledge is used to produce more knowledge of the same quality, it does so at a fixed rate $b > \rho$ per unit of input. In other words the production function for existing knowledge is linear in the knowledge capital used as input. This can be regarded as *capital widening* or competitive imitation.⁶

More knowledge can also be created through innovation – that is, the production of a higher quality of knowledge capital from an existing quality. This can be regarded as *capital deepening*. Specifically, a unit of knowledge capital of quality $j + 1$ can be produced from $a > \lambda$ units of quality j . This represents the conversion of knowledge from, say, radios to TVs. In the strict interpretation of k_j as a stock of knowledge about how to produce, capital widening corresponds to the spread of existing knowledge, and capital deepening to the creation of new knowledge from old.⁷

⁶ As in Grossman-Helpman, the assumption of linearity is purely for algebraic convenience: any sufficiently productive concave function would also do and, in fact, enrich the model from an applied perspective.

⁷ Notice that new knowledge capital loses the capability of the old knowledge capital that was converted. This may be true for the physical replacement of machines with newer models, but is not usually the case for human capital – a pilot may still be capable of flying a Cessna after learning to fly a Stealth bomber. However, the key constraint is that at any moment in time you must be acting as one of the two kinds of pilot, but not as both. If we introduced a technology for converting quality $j + 1$ knowledge capital back to quality j at the same ratio as the forward conversion, then this would precisely capture knowledge that was not lost, but knowledge capital as a resource that could be deployed at only one level at a time. However, in the case we study, there is no reason to use the backward conversion technology, so the equilibrium would not change. Should there be an indivisibility in the deepening technology, the possibility of such a

Alternatively, knowledge capital of quality j can be employed in the production of quality j consumption on a one-to-one basis. As in the Grossman and Helpman model, the production of consumption also requires labor – leading here to diminishing returns for each quality of knowledge capital. Note that it is the diminishing returns on each rung of the ladder that is crucial to our story: the specific details of diminishing returns arising from a labor constraint are used because it is convenient analytically. Specifically, each unit of quality j knowledge capital employed in the consumption sector requires also a unit of labor to produce a unit flow of quality j consumption. Our key assumption is that measured in units of current consumption the creation of new knowledge is costlier than spreading knowledge already in existence, that is $b > \lambda / a$. This implies that as long as it is not needed for expanding consumption – that is, until all labor is employed with the most advanced quality of knowledge capital – it is not socially efficient to introduce a higher quality of knowledge in the production of consumption. As before, we normalize the fixed labor supply to one.

Let h_j denote the flow investment of knowledge capital of quality j in the production of knowledge capital of quality $j + 1$. Under our assumptions the motion over time of quality j stock of knowledge is given by

$$\dot{k}_j + h_j = b(k_j - d_j) + \frac{h_{j-1}}{a}.$$

Notice that d_j units of the stock of knowledge must be allocated to the consumption sector, so the amount available for production of additional knowledge of any quality, is $k_j - d_j$. We require that $d_j \leq k_j$ and $h_j \geq 0$. The flow of new knowledge must be divided into an increase in existing knowledge, \dot{k}_j and investment in higher quality knowledge, h_j . Note that over a short period of time, there is no limit on the flow of quality j knowledge into quality $j + 1$, that is h_j may be arbitrarily large, provided \dot{k}_j is correspondingly negative, so we allow also discrete conversion $\Delta k_{j+1} = -\Delta k_j / a$.

The key technical fact is that this economy is an ordinary diminishing return economy. Competitive equilibrium can be described in the usual way as the combination

backward conversion would effectively eliminate the indivisibility: just produce enough of the new knowledge to overcome the indivisibility, then convert the excess back to the old knowledge. This may help explain why, in practice, indivisibilities do not seem to be an important practical problem for the bulk of innovation.

of consumer optimization and profit maximization. The first and second welfare theorems hold, so efficient allocations can be decentralized as a competitive equilibrium and vice versa.⁸

5. *Competitive Equilibrium with Knowledge Capital*

Our first goal is to characterize the competitive equilibrium of the model. We will examine the robustness of the equilibrium to the main assumptions later. Here we will show that as the stock of knowledge capital grows, after a possible *initial unemployment phase* where there is too little knowledge to employ the entire stock of labor, the competitive equilibrium settles into a recurring cycle.

The cycle alternates between a *growth* phase (widening) in which consumption grows at the rate $b - \rho$ by upgrading the stock of knowledge used to produce it, and a *build-up* phase (deepening) in which consumption remains flat while a new quality of knowledge capital is accumulated. The growth phase ends when it is no longer possible to increase consumption without innovation, that is when all labor is applied to the most advanced type of knowledge capital, say quality j .

To understand how innovation takes place – that is, knowledge capital of quality $j + 1$ is first created, then progressively accumulated and applied to the production of consumption by replacing knowledge capital of quality j – the key idea is that the relative value of quality j and $j + 1$ knowledge capital simultaneously used in producing consumption is λ , while the instantaneous price ratio, when the innovation first occurs, must be $a > \lambda$. This has an important implication. Because $a > \lambda$, at the first moment that the switchover from quality $j - 1$ to quality j knowledge capital is complete using knowledge capital of quality $j + 1$ to replace knowledge capital of quality j in the production of consumption, would yield negative profits. A positive amount of time – the “build up phase” – must elapse between the exhaustion of the $j - 1$ technology and the use of knowledge capital of quality $j + 1$ to produce consumption. During this build up phase the new type of knowledge is accumulated and its price drops

⁸ The welfare theorems are proven for the discrete time version of this model by Boldrin and Levine [2002]. In that paper we assumed that one unit of capital of type j and λ^{-j} units of labor were required to produce a unit of quality j consumption, and that the capital is used up in the process of producing consumption. This simply changes the units in which knowledge capital is measured. The assumptions here are chosen for compatibility with Grossman and Helpman [1991].

until it becomes cheap enough to use in the production of consumption. This begins the next growth phase, which, therefore, comes a discrete amount of time after the end of the previous one.⁹

Research intensity, the inverse of the combined length of the two phases, we will compute to be

$$j^* = \frac{b - \rho}{\log a}.$$

The Pricing of Knowledge Capital

The key to understanding the competitive equilibrium is the pricing of knowledge capital. It is convenient not to take labor as numeraire, but rather to take current utility as numeraire. This implies that the current price of consumption is marginal utility; in our logarithmic case equal to $1/c_t$. Define q_{jt} to be the time t price of quality j knowledge capital. We first examine the value of knowledge in the production of more knowledge.

One use of knowledge is for innovation – that is to create higher quality knowledge capital. *Zero profits on innovation* implies that $q_{j+1,t} - aq_{jt} \leq 0$ or, equivalently, $q_{j+1,t} / q_{jt} = a$ at the time of innovation. Knowledge can also be used for imitation – that is to create more knowledge of the same quality. In this case the physical rate of return is the growth rate b , so we must have this return plus capital gains equal to the subjective interest rate, that is, $b + \dot{q}_{jt} / q_{jt} = \rho$, or equivalently for *zero profits on imitation* $\dot{q}_{jt} / q_{jt} = -(b - \rho) < 0$.¹⁰

That the price of knowledge is necessarily falling over time is significant. Consider, for example, the fact that the first mover must incur greater costs than subsequent competitors. This is true here, since, to produce quality $j + 1$ from quality j , the first innovator must pay more than subsequent imitators for knowledge as an input. However: by virtue of being first, he can also sell his freshly created knowledge for a higher price than his imitators, who must sell at a later date when the new knowledge is

⁹ We are grateful to V.V. Chari for forcing us to clarify the intuitive reason why the build up phase must have a strictly positive length.

¹⁰ Strictly speaking, these conditions need only hold with inequality if knowledge capital is not being used to produce more knowledge capital. For example, if there is no capital of quality J or higher, then we can have $q_{j+1,t} < aq_{jt}$ for $j + 1 \geq J$. However, there exist equivalent equilibria in which profits are zero. See Boldrin and Levine [2002] for the technical details.

worth less. In other words, even in this model of perfect competition, there is a first mover advantage, because the competitive price of knowledge capital is falling over time. It is this first mover advantage that motivates innovators to act while knowing they will soon be imitated.

Consumption Value of Knowledge Capital: The Initial Unemployment Phase

Next we turn to the value and price of knowledge used to produce consumption. When the economy begins, the initial condition may be such that there is insufficient knowledge capital to employ all the labor. We first examine the value of knowledge used to produce consumption during this phase. Recall that, when they are both used in producing consumption, the relative value of quality j and $j + 1$ knowledge is λ , while at the time the second is first created their price ratio must be $a > \lambda$. This implies that, when there is no full employment, only the lowest quality of knowledge is used to produce consumption. Subsequently we will establish that when there is full employment, no more than two qualities of knowledge capital are used to produce consumption.

What is the value of knowledge capital used to produce consumption when there is unemployment? We can determine this from the fact that consumption and the knowledge capital producing it are perfect substitutes: both can and are produced from the current stock of knowledge.¹¹

Proposition 1: *There is an initial unemployment phase during which the single lowest quality of knowledge, j , is used to produce consumption. The price of this knowledge capital is $q_{jt} = v_{jt} = \lambda^j / bc_t$. Consumption grows at the constant rate $\dot{c}_t / c_t = b - \rho$. The unemployment phase ends when $c_t = \lambda^j$.*

Proof: Consider how we might use a small amount ε of quality j knowledge over a short period of time τ to produce either consumption or more knowledge of quality j . If consumption is to be produced, ε units of quality j knowledge yield $\lambda^j \varepsilon \tau$ units of consumption. If quality j knowledge is to be produced the yield is $b \varepsilon \tau$ new units. Hence one unit of quality j knowledge is a perfect substitute for λ^j / b units of consumption.

Since the marginal social value of c_t units of quality adjusted consumption is $1/c_t$, we can conclude that, when there is unemployment, the marginal social value in

¹¹ By assumption the labor constraint is not binding during the initial unemployment phase.

producing consumption of a unit of quality j knowledge is $v_{jt} \equiv \lambda^j / bc_t$. Its price cannot be less than this: we must have $q_{jt} \geq v_{jt}$, with equality if knowledge of that quality is actually used to produce consumption. As claimed, if quality j' knowledge capital is used to produce consumption, then no higher quality can be used. That is, if quality $j > j'$ was used to produce consumption, then we would have $q_{jt} / q_{j't} = v_{jt} / v_{j't} = \lambda^{j-j'}$. But, because knowledge capital of quality j can be readily obtained from that of quality j' by applying the innovation technology for $j - j'$ times, the latter equality would contradict the zero profit on innovation condition $q_{jt} / q_{j't} = a^{j-j'}$.

We have shown that during the *initial unemployment phase* a single quality of knowledge, j , is used to produce consumption, the lowest available quality. The price of this knowledge capital is $q_{jt} = v_{jt} = \lambda^j / bc_t$. Since the no profit on imitation condition is $\dot{q}_{jt} / q_{jt} = -(b - \rho)$, it must be that consumption grows at the constant rate $\dot{c}_t / c_t = b - \rho$. Eventually, the stock of knowledge of quality j grows sufficiently large that a quantity of consumption $c_t = \lambda^j$ can be produced, meaning that we have reached full employment and that it is no longer possible to increase output by employing more quality j knowledge.

☑

Consumption Value of Knowledge Capital: Full Employment

We next examine the value of knowledge capital used to produce consumption when there is full employment. Suppose that quality $j' < j$ knowledge is used to produce consumption. In this case, consumption may be increased by replacing some of the inferior quality j' knowledge capital with additional units of quality j knowledge, which has, therefore, value in producing consumption. Notice that, instead, the lowest quality knowledge used in producing consumption has no marginal value there: additional units cannot increase consumption at all.

We follow the same analysis as in the initial unemployment phase: we determine a price ratio between consumption and quality j knowledge capital by examining how much of each we can produce from a given small amount ε of quality j knowledge over a short period of time τ . The key difference is that with full employment when we move quality j knowledge into the consumption sector, we must displace some lower quality

j' knowledge in order to free up the labor needed to work with the newly added quality j knowledge.¹² This frees the displaced j' knowledge, which, of course, has social value.

Proposition 2: *During full employment no more than two qualities of knowledge capital are actually used to produce consumption, and these must be consecutive qualities. If j' is used to produce consumption*

$$q_{jt} \geq v_{jt}^{j'} \equiv \frac{\lambda^j - \lambda^{j'}}{b(1 - 1/a^{j-j'})} \frac{1}{c_t},$$

with equality if j is also used to produce consumption. If two qualities are both used in positive amounts in producing consumption the latter grows at the rate $b - \rho$.

Proof: The easiest way to do the computation is to imagine that newly freed quality j' knowledge is converted immediately to quality j . Consequently, to add one unit of quality j knowledge to the production of consumption requires drawing less than one unit of quality j knowledge away from other uses.

Specifically, increasing the quantity of quality j knowledge used in the production of consumption by $\varepsilon/(1 - 1/a^{j-j'})$ displaces the identical amount of quality j' knowledge. These $\varepsilon/(1 - 1/a^{j-j'})$ free units of quality j' knowledge convert to

$$1/a^{j-j'} \times \varepsilon/(1 - 1/a^{j-j'})$$

units of quality j knowledge. The net input of quality j knowledge needed to increase the amount used in consumption by $\varepsilon/(1 - 1/a^{j-j'})$ is the difference between the total requirement $\varepsilon/(1 - 1/a^{j-j'})$ and this extra quality j knowledge generated from freed knowledge j' capital

$$\frac{\varepsilon}{1 - 1/a^{j-j'}} - \frac{1}{a^{j-j'}} \frac{\varepsilon}{1 - 1/a^{j-j'}} = \varepsilon.$$

In other words, if we move ε units of quality j knowledge to displace quality j' in the production of consumption, the amount of quality j knowledge used in producing consumption increases by the greater amount $\varepsilon/(1 - 1/a^{j-j'})$. Another way of seeing

¹² This is often labeled “creative destruction”: a more productive technology replaces a less productive one through the scrapping of the latter and the shifting of workers from old plants to new ones. In practice, this often involves replacing firms using the old quality capital with firms that have adopted the new one.

this is: investing ε additional units of quality j knowledge capital in the production of consumption allows to shift $\varepsilon/(1 - 1/a^{j-j'})$ units of labor away from quality j' and toward quality j , the total installed capacity of which increases by $\varepsilon/(1 - 1/a^{j-j'})$.

We conclude that, over a short period of time τ , using an additional ε units of quality j knowledge capital to displace the inferior quality j' in the production of consumption increases consumption by

$$(\lambda^j - \lambda^{j'}) \frac{\varepsilon}{1 - 1/a^{j-j'}} \tau,$$

where $\lambda^j - \lambda^{j'}$ is the productivity differential, measured in consumption units, between the high and low quality knowledge.

Turning next to the use of knowledge in reproducing itself, as was the case in the initial unemployment phase, if quality j knowledge capital is used to reproduce itself, it yields $b\varepsilon\tau$ new units.

Putting together the two uses of knowledge in the full employment case, we conclude that one unit of j knowledge capital is a perfect substitute for

$$\frac{\lambda^j - \lambda^{j'}}{b(1 - 1/a^{j-j'})}$$

units of consumption, when knowledge of type j' is also used to produce consumption. This gives the marginal social value of a unit of quality j knowledge in producing consumption by displacing quality j' knowledge as

$$v_{jt}^{j'} = \frac{\lambda^j - \lambda^{j'}}{b(1 - 1/a^{j-j'})} \frac{1}{c_t}.$$

The price q_{jt} of quality j knowledge capital cannot be lower than $v_{jt}^{j'}$, since it cannot be strictly profitable to buy knowledge of type j and shift it into the production of consumption. So if j' is used to produce consumption

$$q_{jt} \geq v_{jt}^{j'},$$

with equality if j is also used to produce consumption.

Let us now show that only two qualities of knowledge capital are actually used to produce consumption and, when this is the case, consumption grows at the rate $b - \rho$. We make use of the following calculation:

$$\begin{aligned} v_{jt}^{j'} &= \frac{\lambda^j - \lambda^{j'}}{b(1 - 1/a^{j-j'})} \frac{1}{c_t} = a^{j-j'} \lambda^{j'} \frac{\lambda^{j-j'} - 1}{a^{j-j'} - 1} \frac{1}{bc_t} \\ &= a^{j-j'} \lambda^{j'} \frac{\lambda^{j-j'-1} + \lambda^{j-j'-2} \dots + 1}{a^{j-j'-1} + a^{j-j'-2} \dots + 1} \frac{\lambda - 1}{a - 1} \frac{1}{bc_t} \\ &< a^{j-j'} \lambda^{j'} \frac{\lambda - 1}{a - 1} \frac{1}{bc_t} = a^{j-j'-1} \left[a \lambda^{j'} \frac{\lambda - 1}{a - 1} \frac{1}{bc_t} \right] = a^{j-j'-1} v_{j'+1,t}^{j'} \end{aligned}$$

Now suppose that quality j' and quality $j > j'+1$ are both used to produce consumption. Then $q_{jt} = v_{jt}^{j'}$. Zero profits on innovation implies $q_{jt} = a^{j-j'-1} q_{j'+1,t}$. Using $q_{j'+1,t} \geq v_{j'+1,t}^{j'}$, we have

$$a^{j-j'-1} v_{j'+1,t}^{j'} \leq a^{j-j'-1} q_{j'+1,t} = q_{jt} = v_{jt}^{j'} < a^{j-j'-1} v_{j'+1,t}^{j'},$$

the last step following from the inequality above. This contradiction establishes that if j' is used to produce consumption, then no higher quality than $j'+1$ may be used. In other words, at most two adjacent qualities of knowledge capital are used to produce consumption at any given point in time. When this is the case, we must have, for quality j , the price $q_{jt} = v_{jt}^{j-1}$. In particular, $\dot{v}_{jt}^{j-1} / v_{jt}^{j-1} = \dot{q}_{jt} / q_{jt} = -(b - \rho)$. Plugging in for v_{jt}^{j-1} this implies once again that consumption grows at $b - \rho$. □

The Growth Cycle

We are now in a position to analyze what happens at the end of the initial unemployment phase when employment is full. The economy cycles between growth phases during which consumption grows at the constant rate $b - \rho$, and build-up phases where consumption remains constant. After each complete cycle, it repeats at the next level of the quality ladder.

Proposition 3: *Consumption remains constant during a build-up phase that lasts for*

$$\tau^b = \frac{\log a - \log \lambda}{b - \rho}$$

followed by a growth phase during which consumption grows at the rate $b - \rho$ lasting

$$\tau^g = \frac{\log \lambda}{b - \rho}$$

The total length of a cycle is

$$\tau = \frac{\log a}{b - \rho}$$

so that the resulting research intensity is

$$j^* = \frac{b - \rho}{\log a}$$

Proof: We may assume without loss of generality that the initial quality used is $j = 0$. If the unemployment phase ends at t , then $c_t = 1$. The price $q_{0t} = v_{0t} = 1/b$, from which we can see that the price of quality $j = 1$ knowledge is $q_{1t} = a/b$ at the time it is introduced. The value of this knowledge capital in producing consumption is

$$v_{1t}^0 = \frac{\lambda - 1}{b(1 - 1/a)} = \frac{a}{b} \frac{\lambda - 1}{a - 1} < q_{1t}.$$

That is, right when it is first produced, it does not pay to use quality $j = 1$ knowledge to produce consumption. It follows that consumption must now remain constant for some time; as long as consumption remains constant, so does the value of knowledge of quality $j = 1$ in producing it, $v_{1,t+\tau}^0 = v_{1t}^0$. On the other hand, the price of quality $j = 1$ knowledge capital is falling as it accumulates over time, so that, at $t + \tau$ it is $q_{1,t+\tau} = (a/b)e^{-(b-\rho)\tau}$. When

$$v_{1t}^0 = \frac{a}{b} \frac{\lambda - 1}{a - 1} = (a/b)e^{-(b-\rho)\tau} = q_{1t}$$

that is, at

$$\tau = \frac{1}{b - \rho} \log \frac{a - 1}{\lambda - 1}$$

it becomes profitable to introduce quality $j = 1$ knowledge into producing consumption, and consumption then grows at the rate $b - \rho$.

In general, when quality j knowledge capital is first created quality adjusted consumption is λ^{j-1} and further increases require quality j to be used in its production until consumption reaches λ^j . During the growth phase, qualities $j - 1$ and j are used,

and consumption grows at $b - \rho$. Hence the length of the growth phase τ^g is characterized by

$$\lambda^{j-1} e^{(b-\rho)\tau^g} = \lambda^j$$

meaning that

$$\tau^g = \frac{\log \lambda}{b - \rho}.$$

Applying to the general case the argument developed earlier for qualities $j = 0$ and $j = 1$, we see that, at the end of the j -th growth phase, the price of quality $j + 1$ knowledge capital must be

$$q_{j+1,t} = a q_{jt} = a v_{jt}^{j-1} = a \frac{\lambda^j - \lambda^{j-1}}{b(1 - 1/a)} \frac{1}{\lambda^j} = \frac{\lambda - 1}{b(1 - 1/a)} \frac{a}{\lambda}.$$

The consumption value of quality $j + 1$ knowledge capital is

$$v_{j+1,t}^j = \frac{\lambda^{j+1} - \lambda^j}{b(1 - 1/a)} \frac{1}{\lambda^j} = \frac{\lambda - 1}{b(1 - 1/a)},$$

implying, again, that $v_{j+1,t}^j < q_{j+1,t}$. In other words, at the end of each growth phase, we must have a *build-up* phase, during which consumption remains fixed at λ^j , while the price of quality $j + 1$ knowledge (which is being accumulated) falls by a factor of λ/a . Since it falls at the constant rate $b - \rho$, this takes

$$\tau^b = \frac{\log a - \log \lambda}{b - \rho}.$$

Following this, we again begin the growth phase, and the cycle repeats at the next level of the quality ladder.

The total length of the cycle follows from adding $\tau^g + \tau^b$, and the research intensity – the rate at which we move up the ladder – is by definition is just the inverse of the cycle length.

☑

6. Comparison of the Models

We have, now, three possible models explaining endogenous growth. One is the Grossman-Helpman model, in which the innovation rate is given by

$$\iota = \frac{(1 - 1/\lambda)}{a_I} - \frac{\rho}{\lambda}.$$

Another is the Grossman-Helpman efficient solution, and here the innovation rate is given by

$$\iota^* = \frac{1}{a_I} - \frac{\rho}{\log \lambda}.$$

Finally, we have our model of competitive knowledge capital accumulation, in which the innovation rate is given by

$$j^* = \frac{b - \rho}{\log a}.$$

Each model is designed to get a closed form solution. Qualitatively, each is a similar function of the cost of innovating and the degree of impatience. As consumers are more patient, the frequency of innovation goes up. As it becomes more costly to innovate, innovation goes down. There are of course minor differences in the functional forms between these solutions. But the functional forms depend on a variety of assumptions – logarithmic utility, exponentially improving steps, and so forth – that were contrived to make the models easy to solve, so the particular functional forms have no strong claim of correctness. Moreover, the models differ in ways that are designed to ease the solution. For example, it is technically convenient for Grossman and Helpman to assume the same labor is used in producing knowledge as in producing output. However, it is technically convenient for us to assume that knowledge capital is produced only from knowledge capital.

What are then the substantive, as opposed to the technically convenient, differences? First, the parameter λ , how high each rung of the ladder is, has no effect in our model – this is due to the presence of two offsetting effects, increasing the intensity of innovation during the first build-up phase of the cycle – this effect being present also in Grossman-Helpman – and decreasing it during the second growth phase. However, this “neutrality” of step size in our model is due to the simplifications involved, and it is likely that various adjustments could yield either increasing innovation in step-size as in Grossman-Helpman, or decreasing innovation in step-size.


Second, the competitive innovation model has the extra parameter b , representing the rate at which productive capacity increases or is turned into usable output, possibly

through imitation. As it becomes easier to reproduce knowledge capital, in the sense that b is larger, the intensity of innovation increases. In a certain sense the Grossman-Helpman model, like all models of that class, assumes that $b = \infty$. This is because once the fixed cost is paid and a new rung has been introduced the technology allows one to make an infinite number of copies of the j th good: a finite number of copies is made only because the monopoly power of the first innovator is used to prevent competitive imitation. Put it differently, the Grossman-Helpman model assumes that the movement from one particular vintage to the next can be infinitely rapid (once it is discovered knowledge is a public good and everyone has it); the presence of a fixed cost at each step explains why things do not go haywire.

Finally, the model of competitive knowledge capital accumulation has a distinctive implication for the growth of output: it should grow in bursts (the growth phase) punctuated by periods in which output does not grow (the build-up phase). Moreover, the build-up phase should end when a new vintage of consumption is produced for the first time. Notice first that, per-se, the fact that the build-up phase has a positive length is not essential to our theory: when $a = \lambda$ we have that $\tau^b = 0$ and the new knowledge is used to produce consumption right away. Still, as long as $b > \lambda/a$ meaning that innovation is costlier than imitation, the former takes place only when all labor is employed with the latest available quality of knowledge. Further, the growth process still alternates between widening and deepening even if the latter is instantaneous and consumption grows at a constant balanced rate, as in most models of endogenous growth. It should be relatively easy to see that, in this special case, our model “endogenizes” Solow’s growth model of vintage capital with exogenous TFP growth.

Turning to the DRAM data from Irwin and Klenow [1994] we find that new vintages in fact trickle into production which is at variance with the model. This is not terribly surprising, as in practice even during the buildup phase there will be some demand for larger capacity chips for specialty use, and some chips produced for test purposes and subsequently sold for specialty use.

So let us take as our operational definition of “produced for the first time” the first time the new generation constitutes 5% of the total market for memory. We can then examine aggregate output of memory (that is measured in bytes, not chips) and plot this against the dates at which new generations of chips are introduced.



The table above shows the aggregate memory output as the blue curve, with the dates at which new generations are introduced as vertical red lines. The data does not show that every slowdown is followed by a switchover, nor even that every switchover is preceded by a slowdown (the 1984 switchover clearly is not), but rather that, in general, there are period of growth alternating with slowdowns. Remarkably these slowdowns have essentially zero growth as the theory predicts, and they are clearly associated with the switchovers. The ratio between the growth and buildup phases in the data is interesting too, as the buildup phases are very short, suggesting that λ is not much smaller than a .

7. Fixed Cost of Knowledge Capital

There can be little doubt as an empirical fact that there is a fixed cost in creating new knowledge. Of course there are fixed costs in producing just about everything, and this has not prevented the competitive model with perfect divisibility from being a useful tool in studying a wide range of market phenomenon. Never-the-less, it would be reassuring to know that the introduction of fixed costs into our model of competitive innovation does not cause things to collapse.

Let us assume, then, that there is a fixed cost of introducing new knowledge capital for the first time. Specifically, to produce for the first time quality $j + 1$ knowledge from quality j requires a fixed cost of F units of quality j knowledge. This results in the creation of $\bar{k} < F$ initial units of quality $j + 1$ knowledge.¹³ For computational simplicity and notational convenience, once the fixed cost is incurred, we assume that it is possible to convert additional units of quality j knowledge to quality $j + 1$ knowledge at the same rate $F / \bar{k} = a$. This may seem a strong assumption – it might seem plausible that, after knowledge of a given quality is first converted to new knowledge capital, the cost of converting additional old knowledge capital will fall. It will become clear later that, properly modeled, such an assumption adds complication to the model, without changing the essential results.

We will assume throughout that

$$\bar{k} \leq k^* \equiv \frac{a^{\frac{b}{b-\rho}} - 1}{\frac{b}{a^{\frac{b}{b-\rho}} - a}}.$$

To understand this assumption, observe that $k^* > 1$. If $\bar{k} > 1$ then the first productive capacity installed saturates the entire market for knowledge capital. In fact, if \bar{k} is even larger, $\bar{k} > k^*$, then we can show that any attempt to innovate must result in the price of the newly created knowledge falling to zero, meaning that innovation will not take place under competition. This may be thought of as the classical case – the case that Romer, Grossman and Helpman, Aghion and Howitt and other apparently have in mind. It seems

¹³ As discussed in footnote 7, if it is possible to convert new knowledge capital to old at the same ratio as forward conversion, then fixed cost does not matter, as it can be “undone” by a subsequent backward conversion. This may partially explain why, when we consider the empirics of fixed costs, they do not seem to matter a great deal: most electrical engineers, if not most economists, can still change a light bulb by themselves.

implausible to us that the first unit of productive capacity for a new good will saturate the market for that good – that nobody would want even a single additional unit of productive capacity. Of course these subsequent units might cost considerably less than the original. This is the case in our model, since knowledge capital devoted to copying grows at the rate b , and its price falls exponentially over time.

It is convenient to take as parameters F, a . This enables us to compare the indivisible case more directly with the divisible case. Implicitly it means that we vary \bar{k} as F changes. So our assumption on \bar{k} means that

$$F \leq \frac{a^{\frac{b}{b-\rho}} - 1}{a^{\frac{\rho}{b-\rho}} - 1}.$$

In the perfectly divisible model, the exact time at which quality j knowledge is converted to quality $j+1$ knowledge is a matter of indifference. Any quality j knowledge not being used to produce consumption may equally well be used to produce more knowledge of the same quality, to be converted to a higher quality at a later date, or it can equally well be converted right now. Amongst all these equilibria, there is one at which knowledge of quality j is not converted to knowledge of quality $j+1$ until the moment at which quality $j+1$ knowledge is used for the first time in the production of consumption – that is, at the end of build-up and beginning of growth. Moreover, there is a unique such equilibrium in which all quality j knowledge not being used to produce consumption is converted to quality $j+1$ at that time. Along the equilibrium path, let F^* denote the unique amount of quality j knowledge not being used in the production of consumption at the end of build-up, and let t be the time at which that build up ends. Then $F^* = k_{j, \tau^g + \tau^b} - 1$. Since the equilibrium is a steady state cycle, this is independent of j , and in the Appendix we show that

$$F^* = \frac{a^{\frac{\rho}{b-\rho}}}{a^{\frac{\rho}{b-\rho}} - 1} \frac{\lambda^{\frac{\rho}{b-\rho}} - 1}{\lambda^{\frac{\rho}{b-\rho}}} \frac{a-1}{\lambda-1} \frac{b-\rho}{\rho}.$$

In other words: along the equilibrium cycle of our model, even when a technological indivisibility is not present, there is a moment – when the build up ends and growth resumes - at which a large amount of new productive capacity is installed all at once. An external observer could interpret such large initial investment in productive

capacity as due to the presence of a technological fixed cost in the innovative activity but, as we have seen, this needs not be the case. Clearly, when a true fixed cost of innovation is assumed, two cases are possible: the case of small fixed cost in which $F \leq F^*$, and the case of large fixed cost in which $F > F^*$.

The case of small fixed cost is easy to handle: there is an equilibrium of the divisible economy in which the constraint of fixed cost does not bind. That is, there is an equilibrium in which no one chooses to convert quality j knowledge to quality $j + 1$ until the end of the build-up, at which time the amount converted is greater than the fixed cost of innovation. Put differently, the fixed cost constraint – that knowledge must be converted in a minimal amount – does not bind. This gives us a minimal robustness property – if the fixed cost is small enough, nothing changes, and competition is free to work its magic.

In the case of large fixed cost no competitive equilibrium exists. Notice that this is a different statement than “no innovation will take place due to fixed costs,” or “competition is impossible.” It is simply a statement about a particular model of competitive equilibrium. It is possible to introduce an underlying game-theoretic model of competitive traders and work out the equilibrium of such a game, but that must await further research.

8. Robustness

We have developed a model based on a number of special assumptions designed for ease of exposition and solution. We now verify that the basic conclusion – that innovation can take place under competition – is not sensitive to them.

Features of the innovation process

First is the fact that when new knowledge capital is created from old, less of the new knowledge capital is made available. In other words, an additional unit of new knowledge requires more than one unit of the old knowledge. This is a natural assumption: the stock of knowledge capital is measured by the amount of labor it can employ. If the creation of new knowledge capital, for example, is in the form of upgrading old machines to a new technology, there will generally be some wastage in the

conversion process. But more broadly, the fact that when labor embodying old knowledge capital is converted into labor embodying the new knowledge capital, this “new” labor will not have the benefit of time and experience that had accrued to the old. For example, if the knowledge capital is human capital, then the expertise of engineers and others in the new knowledge capital will naturally be less than in the old.

One question is why we assume that new knowledge capital loses the productive capacity of the old. That is, simply because a new skill is acquired does not mean that the old one is lost – although for physical capital it may mean exactly that. However, the critical assumption is that only one of the two skills can be employed at a moment of time by a given individual, not that the old skill is lost: while you may still be capable of flying a Cessna after learning to fly a Stealth bomber, at any given moment in time you must be acting as one of the two kinds of pilot, but not as both. If we introduced a technology for converting quality $j + 1$ knowledge capital back to quality j at the same ratio as the forward conversion, then this would precisely capture knowledge that was not lost, but knowledge capital as a resource that could be deployed at only one level at a time. However, there would be no reason to ever use the backward conversion technology, so the equilibrium consumption path would not change. Note however, that the possibility of converting back at the same ratio as forward would have an important implication in the next section where we consider a fixed cost of conversion. The implication is that the fixed cost would not matter, as it could be “undone” by a subsequent backward conversion. This may partially explain why, when we consider the empirics of fixed costs, they do not seem to matter a great deal.

Finally, there is an aspect of the conversion technology that may strike the reader as unrealistic: the conversion takes place instantly – while in fact the development of new knowledge generally takes time. This however turns out to be a matter of convenience rather than substance. Suppose that a units of quality j knowledge capital converted at time t do not become available as a unit of quality $j + 1$ knowledge capital until time $t + \Delta$. Since the model is one of perfect foresight, the need for conversion can be foreseen, so if we imagine the conversion taking place at time $t + \Delta$ then the a units of quality j knowledge capital not being used to produce consumption would have grown to $ae^{b\Delta}$. Hence if we consider a model with no delay but cost of conversion $a' = ae^{b\Delta}$,

this is equivalent to the model with delay and cost of conversion a . In effect, all we have done is to assume that any time delay is capitalized into the cost of conversion.

Labor Cost of Producing Knowledge

We have assumed, *contra* Grossman-Helpman, that labor is not an input into the knowledge creation process. On the one hand, we doubt that the type of labor used in knowledge creation is a particularly good substitute for the labor used to produce output, so we do not view the Grossman-Helpman assumption as especially realistic either. On the other hand: what happens if we require some sort of labor or other input into the knowledge creation process? In the divisible case, this obviously does not matter: the welfare theorems hold regardless of the details of the production process. In the fixed cost case, this creates an incentive to innovate gradually so as to avoid dragging labor out of the production of consumption all at once. Still, the sensible model would be one in which new knowledge accumulates gradually, but is not useful until a threshold is crossed. That is, the right place to put the constraint is in \bar{k} and not in F . This just means that knowledge capital is not employed in producing output until after \bar{k} units are acquired – either way, our analysis would not change in an important way.

Spillover Externalities

With a spillover externality the producer of knowledge capital does not own all the knowledge capital that she produces: some of it spills over onto other traders who get it for “free.” In the extreme it all spills over, and the original producer gets nothing, but this is not a terribly interesting model. In practice the original innovator gets some share of the knowledge capital produced, share which may be larger or smaller depending on circumstances: no doubt in practice there is always some spillover. This will delay innovation even in the perfectly divisible case, as the innovator will receive only a portion of the marginal social benefit of her innovation rather than all of it. The resulting equilibrium will not be first best of course; however, innovation will still take place, it will simply take longer before the price of inputs drops sufficiently to make it profitable to innovate. As there seems to be little or no empirical evidence that spillover externalities are substantial in practice, the departure from the first best is likely quite small.

Cost of Converting Old Knowledge

We have assumed that after the initial fixed cost is paid the cost of converting further units of old knowledge into new remains unchanged. If we assume that the cost of conversion drops discretely after the initial unit of knowledge capital is produced, this is equivalent to a 100% spillover externality and no innovation will take place. However, that assumption makes little economic sense. We can distinguish two techniques for learning how to convert old knowledge into new: learning from scratch, and being taught. The former does not get easier just because somebody already did it. The latter requires the new knowledge as input. This suggests that the correct model is one with an additional activity: one that uses new knowledge to convert old knowledge to new. Consider first the divisible case. Intuitively the social optimum is to create a small amount of new knowledge using the learning from scratch technique, then switch over to the conversion activity. The key point is that, in this case as in our simpler model, it is not socially optimal to create an infinitesimal amount of new knowledge from scratch and then switch over right away. This would provide very little capacity for conversion. Rather there is some threshold in the amount of learning by scratch that is desirable before switching. This, then, is similar to our earlier analysis: the divisible case satisfies the welfare theorems as it simply has an additional activity. Since the divisible equilibrium involves accumulating a decent amount of knowledge using the “from scratch” activity before switching to the “conversion” activity, an indivisibility requiring a certain amount of “from scratch” learning need not bind.¹⁴

9. Conclusion

In our account of innovation, fixed costs and monopoly power play a peripheral role. In the traditional account they are at center stage. This is because, we imagine, of a sequence of misunderstandings: that accounting for endogenous innovation and growth requires increasing returns to scale; that the source of increasing returns to scale is fixed cost; and that, without monopoly, fixed costs are an insurmountable obstacle. For this reason, little evidence about the importance of fixed costs is ever examined. Yet fixed costs and the indivisibilities that give rise to them are ubiquitous. Virtually every type of

¹⁴ The interested reader should consult Boldrin and Levine [2009b], where a model of discovery with the required properties is developed and analyzed.

manufacturing industry has a minimum plant size; human beings come in single indivisible units, and even such humble commodities as shoes can not be produced as less than a single pair, nor wheat less than a single grain. Yet we acknowledge that most of these indivisibilities are not terribly important.

If fixed costs are potentially unimportant even in the traditional account of innovation, they are much less so in our account of innovation. In our model a small fixed cost is irrelevant completely – and even a large fixed cost may have no impact on the intensity of innovation¹⁵.

It is sensible to ask, then, how important are fixed costs in practice? To the extent that they are small, we may reasonably suppose that they are no more central to understanding innovation, than, say to the study of the production of shoes. As one example, we might cite the automobile industry. From Business Week we learn the basic facts about this industry. First, the cost of a new model is quite large: between \$400 and \$900 million – on the same order as the cost of a major new drug and much more than, for example, the cost of a high budget movie. There are still 14 major automobile producers world-wide – there used to be many more when the automobile used to be an innovation – and in the major markets, the United States especially, they compete on a relatively level footing. Patent and intellectual property protection for automobile designs is essentially non-existent. Once introduced, models may be copied with a lag of about a year, yet once in production, major changes occur roughly only every 4-5 years. So, based on the standard model, we might expect not to see new models introduced, yet we do.

Consider next the example of DRAM with which we started. From Park [2002] we learn: “Typically, each generation of DRAM [is] introduced by a leading firm, induces (sequential) entries of up to 20 producers...Usually, the firms, which succeed in the innovation...come up with various different physical designs...” In other words, intellectual property does not play much of a role in this industry and there are a lot of producers. There is also a substantial fixed cost of building new “fabs” in which to produce chips. What happens in this market is that the first-movers get to sell early in the market where output is low and price high. This is in fact a competitive rent, and not due

¹⁵ We refer to Boldrin and Levine [2009a] for a justification of this statement.

to monopoly at all: output is low and price high because capacity is low and marginal utility is high, as in our model. In particular, a producer would be foolish to artificially hold down production in the early periods where demand is elastic, and everyone is racing to beat them out.¹⁶

Yet more evidence for the relative unimportance of short-term monopoly can be found in the common practice of patent pools¹⁷. For example, the steel patent pool was formed not to inhibit innovation, but because innovation was at a standstill. On top of this, there is almost total lack of evidence that patents have any effect on the rate of innovation.¹⁸

We should mention that one of the purposes of the Grossman-Helpman paper was to show how a model of a quality ladder is equivalent to one in which new goods are introduced, even as old ones continue to be used. The model of product diversity is a perfectly sensible account of the benefit of innovation – but it does not make any important difference to the main story. If there are incentives to innovate when a technology has played out, whether it will continue to be used after innovation is not crucial. In some cases the model of technological replacement is the relevant one. For example computer hardware since the Macintosh has had the same basic elements of screen, mouse, removable storage, cpu, and a disk drive. Yet each has undergone numerous upgrades. The flat screen replaced the CRT, the recordable cd replaced the floppy. On the other hand, the laptop has supplemented but not replaced the desktop. In software the so called “killer apps” have tended to supplement but do not replace each other: the spreadsheet, was not replaced by the word processor, which was not replaced by email, which was not replaced by the web browser, which was not replaced by p2p networking, which was not replaced by instant messaging, which in turn was not replaced by Google. On the other hand, within categories innovation proceeds via upgrades. In spreadsheets: Visicalc was replaced by Lotus 123, and then Excel. In word processing Wordstar was replaced by Wordperfect, then Word. But the bottom line is the question of when people start looking for and talking about the “next killer app.” In the software

¹⁶ See Boldrin and Levine [2004] for a detailed analysis.

¹⁷ See, for example, Shapiro [2001].

¹⁸ Lerner [2002].

industry, as in our model, this happens exactly when the last killer app is not such a killer any more.

Appendix

Consider a growth phase beginning at time 0. Suppose that initial knowledge capital used in consumption is entirely of quality j , and that there are k_{j0} units of that quality. Suppose that consumption starts at c_0 , possibly through an initial jump. It will be convenient to define this jump as $\xi_0 = c_0 / \lambda^j$, with $1 < \xi_0 < \lambda$. At time t consumption has grown to $c_t = c_0 e^{(b-\rho)t}$ until it reaches λ^{j+1} ending the growth phase at time

$$t = \tau^g = \frac{\log(\lambda) - \log(\xi_0)}{b - \rho}.$$

During this time we assume that quality j knowledge capital is converted to quality $j+1$ as soon as it is freed from use in producing consumption. Specifically,

$$\begin{aligned} k_{jt} = d_{jt} &= \frac{\lambda^{j+1} - c_t}{\lambda^{j+1} - \lambda^j} = \frac{\lambda^{j+1} - \lambda^j \xi_0 e^{(b-\rho)t}}{\lambda^{j+1} - \lambda^j} \\ &= \frac{\lambda - \xi_0 e^{(b-\rho)t}}{\lambda - 1}. \end{aligned}$$

So

$$\begin{aligned} h_{jt} &= -\dot{k}_{jt} = -\dot{d}_{jt} \\ &= \frac{(\dot{c}_t / \lambda^j)}{\lambda - 1} = (b - \rho) \frac{(c_t / \lambda^j)}{\lambda - 1} \\ &= \frac{\xi_0 (b - \rho)}{\lambda - 1} e^{(b-\rho)t} \end{aligned}$$

We find then that quality $j+1$ knowledge capital grows according to

$$\begin{aligned} \dot{k}_{j+1,t} &= b(k_{j+1,t} + d_{jt} - 1) + \frac{\xi_0 (b - \rho)}{a(\lambda - 1)} e^{(b-\rho)t} \\ &= b(k_{j+1,t} + \frac{1 - \xi_0 e^{(b-\rho)t}}{\lambda - 1}) + \frac{\xi_0 (b - \rho)}{a(\lambda - 1)} e^{(b-\rho)t} \\ &= b(k_{j+1,t} + \frac{1}{\lambda - 1}) + \frac{(1 - a)b - \rho}{a} e^{(b-\rho)t} \frac{\xi_0}{\lambda - 1} \end{aligned}$$

until τ^g .

The build-up phase comes next and lasts until $t = \tau^g + \tau^b$, at which time quality $j+1$ knowledge is $k_{j+1,t} = 1 + e^{b\tau^b} (k_{j+1,\tau^g} - 1)$. Since $\tau^g + \tau^b = \log a / (b - \rho)$, we have, from above, that

$$\tau^b = \frac{\log(a) - \log(\lambda) + \log(\xi_0)}{b - \rho}, \text{ or}$$

$$k_{j+1, \tau^g + \tau^b} = 1 + \left(\frac{a\xi_0}{\lambda} \right)^{\frac{b}{b-\rho}} (k_{j+1, \tau^g} - 1).$$

We next guess a solution to

$$\dot{k}_{j+1, t} = b(k_{j+1, t} + \frac{1}{\lambda - 1}) + \frac{(1 - a)b - \rho}{a} e^{(b-\rho)t} \frac{\xi_0}{\lambda - 1}$$

of the form $k_{j+1, t} = a_0 + a_1 e^{bt} + a_2 e^{(b-\rho)t}$. By differentiating

$$\dot{k}_{j+1, t} = ba_1 e^{bt} + (b - \rho)a_2 e^{(b-\rho)t},$$

substituting into the differential equation gives

$$\begin{aligned} \dot{k}_{j+1, t} &= b(a_0 + a_1 e^{bt} + a_2 e^{(b-\rho)t} + \frac{1}{\lambda - 1}) + \frac{(1 - a)b - \rho}{a} e^{(b-\rho)t} \frac{\xi_0}{\lambda - 1} \\ &= b \left(a_0 + \frac{1}{\lambda - 1} \right) + ba_1 e^{bt} + \left(ba_2 + \frac{(1 - a)b - \rho}{a} \frac{\xi_0}{\lambda - 1} \right) e^{(b-\rho)t}. \end{aligned}$$

We conclude that

$$a_0 + \frac{1}{\lambda - 1} = 0$$

and

$$(b - \rho)a_2 = ba_2 + \frac{(1 - a)b - \rho}{a} \frac{\xi_0}{\lambda - 1}.$$

This gives

$$a_0 = -1/(\lambda - 1)$$

and

$$a_2 = \frac{(a - 1)b + \rho}{a\rho} \frac{\xi_0}{\lambda - 1}.$$

Plugging back into our guessed solution $k_{j+1, t} = a_0 + a_1 e^{bt} + a_2 e^{(b-\rho)t}$

$$k_{j+1, t} = -\frac{1}{\lambda - 1} + a_1 e^{bt} + \frac{(a - 1)b + \rho}{a\rho} \frac{\xi_0}{\lambda - 1} e^{(b-\rho)t}.$$

The initial condition is

$$k_{j+1,0} = \frac{k_{j0} - d_{j0}}{a} = \frac{k_{j0} - \frac{\lambda - \xi_0}{\lambda - 1}}{a},$$

hence

$$\frac{k_{j0} - \frac{\lambda - \xi_0}{\lambda - 1}}{a} = -\frac{1}{\lambda - 1} + a_1 + \frac{(a-1)b + \rho}{a\rho} \frac{\xi_0}{\lambda - 1},$$

enabling us to find the missing coefficient

$$a_1 = \frac{k_{j0} - \frac{\lambda - \xi_0}{\lambda - 1}}{a} + \frac{1}{\lambda - 1} + \frac{(1-a)b - \rho}{\rho a} \frac{\xi_0}{\lambda - 1}.$$

Consequently

$$k_{j+1,t} = -\frac{1}{\lambda - 1} + \left(\frac{k_{j0} - \frac{\lambda - \xi_0}{\lambda - 1}}{a} + \frac{1}{\lambda - 1} + \frac{(1-a)b - \rho}{\rho a} \frac{\xi_0}{\lambda - 1} \right) e^{bt} + \frac{(a-1)b + \rho}{a\rho} \frac{\xi_0}{\lambda - 1} e^{(b-\rho)t}$$

Recall that the growth phase lasts for

$$\tau^g = \frac{\log(\lambda) - \log(\xi_0)}{b - \rho},$$

so $e^{(b-\rho)\tau^g} = \lambda / \xi_0$ and

$$e^{b\tau^g} = \left(\frac{\lambda}{\xi_0} \right)^{\frac{b}{b-\rho}}.$$

Plugging in we find

$$\begin{aligned} k_{j+1,\tau^g} &= -\frac{1}{\lambda - 1} + \left(\frac{k_{j0} - \frac{\lambda - \xi_0}{\lambda - 1}}{a} + \frac{1}{\lambda - 1} + \frac{(1-a)b - \rho}{\rho a} \frac{\xi_0}{\lambda - 1} \right) \left(\frac{\lambda}{\xi_0} \right)^{\frac{b}{b-\rho}} \\ &\quad + \frac{(a-1)b + \rho}{\rho a} \frac{\xi_0}{\lambda - 1} \frac{\lambda}{\xi_0} \\ &= -\frac{1}{\lambda - 1} + \left(\frac{k_{j0} - \frac{\lambda - \xi_0}{\lambda - 1}}{a} + \frac{1}{\lambda - 1} + \frac{(1-a)b - \rho}{\rho a} \frac{\xi_0}{\lambda - 1} \right) \left(\frac{\lambda}{\xi_0} \right)^{\frac{b}{b-\rho}} \\ &\quad + \frac{(a-1)b + \rho}{\rho a} \frac{\lambda}{\lambda - 1} \end{aligned}$$

while, from above, the terminal stock of knowledge capital is

$$\begin{aligned}
k_{j+1, \tau^g + \tau^b} &= 1 + \left(\frac{a\xi_0}{\lambda} \right)^{\frac{b}{b-\rho}} (k_{j+1, \tau^g} - 1) \\
&= 1 + \left(\frac{a\xi_0}{\lambda} \right)^{\frac{b}{b-\rho}} \times \\
&\quad \left\{ -\frac{1}{\lambda-1} + \left(\frac{k_{j0} - \frac{\lambda - \xi_0}{\lambda-1}}{a} + \frac{1}{\lambda-1} + \frac{(1-a)b - \rho}{\rho a} \frac{\xi_0}{\lambda-1} \right) \left(\frac{\lambda}{\xi_0} \right)^{\frac{b}{b-\rho}} \right. \\
&\quad \left. + \frac{(a-1)b + \rho}{\rho a} \frac{\lambda}{\lambda-1} - 1 \right\} \\
&= 1 + \left(\frac{a\xi_0}{\lambda} \right)^{\frac{b}{b-\rho}} \times \\
&\quad \left\{ -\frac{\lambda}{\lambda-1} + \left(\frac{k_{j0} - \frac{\lambda - \xi_0}{\lambda-1}}{a} + \frac{1}{\lambda-1} + \frac{(1-a)b - \rho}{\rho a} \frac{\xi_0}{\lambda-1} \right) \left(\frac{\lambda}{\xi_0} \right)^{\frac{b}{b-\rho}} \right. \\
&\quad \left. + \frac{(a-1)b + \rho}{\rho a} \frac{\lambda}{\lambda-1} \right\},
\end{aligned}$$

$$\begin{aligned}
k_{j+1, \tau^b} &= 1 + \\
&\quad -\frac{\lambda}{\lambda-1} \left(\frac{a\xi_0}{\lambda} \right)^{\frac{b}{b-\rho}} + \left(\frac{k_{j0} - \frac{\lambda - \xi_0}{\lambda-1}}{a} + \frac{1}{\lambda-1} + \frac{(1-a)b - \rho}{\rho a} \frac{\xi_0}{\lambda-1} \right) a^{\frac{b}{b-\rho}} \\
&\quad + \frac{(a-1)b + \rho}{\rho a} \frac{\lambda}{\lambda-1} \left(\frac{a\xi_0}{\lambda} \right)^{\frac{b}{b-\rho}} \\
&= 1 + \frac{1}{\lambda-1} a^{\frac{b}{b-\rho}} - \frac{\lambda}{\lambda-1} a^{\frac{\rho}{b-\rho}} \\
&\quad - \frac{\lambda}{\lambda-1} \left(\frac{a\xi_0}{\lambda} \right)^{\frac{b}{b-\rho}} + \left(\frac{\xi_0}{\lambda-1} + \frac{(1-a)b - \rho}{\rho} \frac{\xi_0}{\lambda-1} \right) a^{\frac{\rho}{b-\rho}} \\
&\quad + \frac{(a-1)b + \rho}{\rho a} \frac{\lambda}{\lambda-1} \left(\frac{a\xi_0}{\lambda} \right)^{\frac{b}{b-\rho}} \\
&\quad + a^{\frac{\rho}{b-\rho}} k_{j0}
\end{aligned}$$

$$\begin{aligned}
k_{j+1,\tau^b} &= 1 + \frac{1}{\lambda - 1} \left(a^{\frac{b}{b-\rho}} - \lambda a^{\frac{\rho}{b-\rho}} \right) \\
&\quad - \left(1 + \frac{(1-a)b - \rho}{\rho a} \right) \frac{\lambda}{\lambda - 1} \left(\frac{a\xi_0}{\lambda} \right)^{\frac{b}{b-\rho}} \\
&\quad + \left(1 + \frac{(1-a)b - \rho}{\rho} \right) \frac{a^{\frac{\rho}{b-\rho}}}{\lambda - 1} \xi_0 \\
&\quad + a^{\frac{\rho}{b-\rho}} k_{j0} \\
&= 1 + \frac{1}{\lambda - 1} \left(a^{\frac{b}{b-\rho}} - \lambda a^{\frac{\rho}{b-\rho}} \right) \\
&\quad + \frac{(a-1)(b-\rho)}{\rho} \frac{\lambda a^{\frac{\rho}{b-\rho}}}{\lambda - 1} \left(\frac{\xi_0}{\lambda} \right)^{\frac{b}{b-\rho}} + \frac{(1-a)b}{\rho} \frac{a^{\frac{\rho}{b-\rho}}}{\lambda - 1} \xi_0 \\
&\quad + a^{\frac{\rho}{b-\rho}} k_{j0}
\end{aligned}$$

$$\begin{aligned}
k_{j+1,\tau^b} &= 1 + \frac{1}{\lambda - 1} \left(a^{\frac{b}{b-\rho}} - \lambda a^{\frac{\rho}{b-\rho}} \right) \\
&\quad + \frac{a-1}{\rho} \frac{a^{\frac{\rho}{b-\rho}}}{\lambda - 1} \left((b-\rho)\lambda \left(\frac{\xi_0}{\lambda} \right)^{\frac{b}{b-\rho}} - b\xi_0 \right) \\
&\quad + a^{\frac{\rho}{b-\rho}} k_{j0}
\end{aligned}$$

So the steady state condition is $k_{j+1,t^b} = k_{j0}$, or

$$\begin{aligned}
\left(a^{\frac{\rho}{b-\rho}} - 1 \right) k_{j0} &= -1 - \frac{1}{\lambda - 1} \left(a^{\frac{b}{b-\rho}} - \lambda a^{\frac{\rho}{b-\rho}} \right) \\
&\quad - \frac{a-1}{\rho} \frac{a^{\frac{\rho}{b-\rho}}}{\lambda - 1} \left((b-\rho)\lambda \left(\frac{\xi_0}{\lambda} \right)^{\frac{b}{b-\rho}} - b\xi_0 \right)
\end{aligned}$$

Note that this function is increasing in ξ_0 , specifically,

$$\begin{aligned}
\left(a^{\frac{\rho}{b-\rho}} - 1\right) \frac{dk_{j0}}{d\xi_0} &= -\frac{a-1}{\rho} \frac{a^{\frac{\rho}{b-\rho}}}{\lambda-1} \left(b\lambda^{-\frac{\rho}{b-\rho}} \xi_0^{\frac{\rho}{b-\rho}} - b\right) \\
&= -\frac{a-1}{\rho} \frac{a^{\frac{\rho}{b-\rho}}}{\lambda-1} b \left(\left(\frac{\xi_0}{\lambda}\right)^{\frac{\rho}{b-\rho}} - 1\right)
\end{aligned}$$

It follows that

$$F^* = k_{j0} - d_{j0} = k_{j0} - \frac{\lambda - \xi_0}{\lambda - 1} = k_{j0} - \frac{\lambda}{\lambda - 1} + \frac{\xi_0}{\lambda - 1}$$

is also increasing in ξ_0 . For reference we can write out the function

$$\begin{aligned}
F^* = & \frac{-1 - \frac{1}{\lambda-1} \left(a^{\frac{b}{b-\rho}} - \lambda a^{\frac{\rho}{b-\rho}}\right) - \frac{a-1}{\rho} \frac{a^{\frac{\rho}{b-\rho}}}{\lambda-1} \left((b-\rho)\lambda \left(\frac{\xi_0}{\lambda}\right)^{\frac{b}{b-\rho}} - b\tilde{c}_0\right)}{a^{\frac{\rho}{b-\rho}} - 1} - \frac{\lambda - \xi_0}{\lambda - 1}
\end{aligned}$$

The sanity check is when $\xi_0 = \lambda$, meaning that there is no growth phase

$$\begin{aligned}
\left(a^{\frac{\rho}{b-\rho}} - 1\right) k_{j0} &= -1 - \frac{1}{\lambda-1} \left(a^{\frac{b}{b-\rho}} - \lambda a^{\frac{\rho}{b-\rho}}\right) \\
&\quad - \frac{a-1}{\rho} \frac{a^{\frac{\rho}{b-\rho}}}{\lambda-1} ((b-\rho)\lambda - b\lambda) \\
&= -1 - \frac{1}{\lambda-1} \left(a^{\frac{b}{b-\rho}} - \lambda a^{\frac{\rho}{b-\rho}}\right) \\
&\quad + (a-1) \frac{\lambda a^{\frac{\rho}{b-\rho}}}{\lambda-1} \\
&= -1 - \frac{1}{\lambda-1} \left(a^{\frac{b}{b-\rho}} - \lambda a^{\frac{\rho}{b-\rho}}\right) \\
&\quad + \frac{\lambda a^{\frac{b}{b-\rho}}}{\lambda-1} - \frac{\lambda a^{\frac{\rho}{b-\rho}}}{\lambda-1}
\end{aligned}$$

$$\begin{aligned}
\left(a^{\frac{\rho}{b-\rho}} - 1\right) k_{j0} &= -1 - \frac{1}{\lambda - 1} \left(a^{\frac{b}{b-\rho}} - \lambda a^{\frac{\rho}{b-\rho}}\right) \\
&\quad + \frac{\lambda a^{\frac{b}{b-\rho}}}{\lambda - 1} - \frac{\lambda a^{\frac{\rho}{b-\rho}}}{\lambda - 1} \\
&= -1 - \frac{a^{\frac{b}{b-\rho}}}{\lambda - 1} + \frac{\lambda a^{\frac{b}{b-\rho}}}{\lambda - 1} \\
&= a^{\frac{b}{b-\rho}} - 1
\end{aligned}$$

while the direct calculation is

$$\begin{aligned}
1 + ((k_{j0} / a) - 1) e^{bt^b} &= k_{j0} \\
1 + ((k_{j0} / a) - 1) a^{\frac{b}{b-\rho}} &= k_{j0} \\
k_{j0} \left(a^{\frac{\rho}{b-\rho}} - 1\right) &= a^{\frac{b}{b-\rho}} - 1
\end{aligned}$$

Finally, when $\xi_0 = 1$, that is, when there is no jump in consumption at the start of growth

$$\begin{aligned}
F^* &= \\
&\frac{-1 - \frac{1}{\lambda - 1} \left(a^{\frac{b}{b-\rho}} - \lambda a^{\frac{\rho}{b-\rho}}\right) - \frac{a-1}{\rho} \frac{a^{\frac{\rho}{b-\rho}}}{\lambda - 1} \left((b-\rho)\lambda^{-\frac{\rho}{b-\rho}} - b\right)}{a^{\frac{\rho}{b-\rho}} - 1} - 1 = \\
&\frac{a^{\frac{\rho}{b-\rho}} \frac{1}{\lambda - 1} \left(-(a-\lambda) - \frac{a-1}{\rho} \left((b-\rho)\lambda^{-\frac{\rho}{b-\rho}} - b\right)\right) - 1}{a^{\frac{\rho}{b-\rho}} - 1} - 1
\end{aligned}$$

$F^* =$

$$\begin{aligned}
& \frac{a^{\frac{\rho}{b-\rho}} \left[1 + \frac{1}{\lambda-1} \left(-(a-\lambda) - \frac{a-1}{\rho} \left((b-\rho)\lambda^{-\frac{\rho}{b-\rho}} - b \right) \right) - 1 \right] - 1}{a^{\frac{\rho}{b-\rho}} - 1} - 1 = \\
& \frac{a^{\frac{\rho}{b-\rho}} \left[1 + \frac{1}{\lambda-1} \left(-(a-\lambda) - \frac{a-1}{\rho} \left((b-\rho)\lambda^{-\frac{\rho}{b-\rho}} - b \right) - \lambda + 1 \right) \right] - 1}{a^{\frac{\rho}{b-\rho}} - 1} - 1 = \\
& \frac{a^{\frac{\rho}{b-\rho}} \left[1 + \frac{1}{\lambda-1} \left(-(a-1) - \frac{a-1}{\rho} \left((b-\rho)\lambda^{-\frac{\rho}{b-\rho}} - b \right) \right) \right] - 1}{a^{\frac{\rho}{b-\rho}} - 1} - 1 = \\
& \frac{a^{\frac{\rho}{b-\rho}} \left[1 + \frac{a-1}{\lambda-1} \left(-\frac{1}{\rho} \left((b-\rho)\lambda^{-\frac{\rho}{b-\rho}} - b + \rho \right) \right) \right] - 1}{a^{\frac{\rho}{b-\rho}} - 1} - 1 =
\end{aligned}$$

$F^* =$

$$\begin{aligned}
& \frac{a^{\frac{\rho}{b-\rho}} \left[1 + \frac{a-1}{\lambda-1} \frac{b-\rho}{\rho} \left(1 - \lambda^{-\frac{\rho}{b-\rho}} \right) \right] - 1}{a^{\frac{\rho}{b-\rho}} - 1} - 1 = \\
& \frac{a^{\frac{\rho}{b-\rho}}}{a^{\frac{\rho}{b-\rho}} - 1} \frac{\lambda^{\frac{\rho}{b-\rho}} - 1}{\lambda^{\frac{\rho}{b-\rho}}} \frac{a-1}{\lambda-1} \frac{b-\rho}{\rho}
\end{aligned}$$

which is strictly positive.

References

- Acemoglu, Daron and Fabrizio Zilibotti [1997], “Was Prometheus Unbound by Chance? Risk, Diversification, and Growth,” *The Journal of Political Economy*, **105**, 709-751.
- Aghion, Philippe and Peter Howitt [1992], “A Model of Growth Through Creative Destruction,” *Econometrica*, **60**, 323-361.
- Andolfatto, David and McDonald, Glenn [1998], “Technology Diffusion and Aggregate Dynamics,” *Review of Economic Dynamics*,

- Boldrin, Michele and David K. Levine [2001], "Growth Cycles and Market Crashes," *Journal of Economic Theory*, **96**:13-39.
- Boldrin, Michele and David K. Levine [2002], "Factor Saving Innovation," *Journal of Economic Theory*, **105**: 18-41.
- Boldrin, Michele and David K. Levine [2004], "Rent Seeking and Innovation," *Journal of Monetary Economics*, **51**: 127-160.
- Boldrin, Michele and David K. Levine [1997, 2008a], "Perfectly Competitive Innovation," *Journal of Monetary Economics*, **55**: 435-453.
- Boldrin, Michele and David K. Levine [2008b], *Against Intellectual Monopoly*, Cambridge University Press.
- Boldrin, Michele and David K. Levine [2009a], "Competitive Entrepreneurial Equilibrium", manuscript, Washington University in St Louis, January.
- Boldrin, Michele and David K. Levine [2009b], "A Model of Discovery", manuscript, Washington University in St Louis, January; forthcoming in *The American Economic Review Papers and Proceedings*, **99** (2), May.
- http://www.businessweek.com/magazine/content/02_45/b3807117.htm
- Funk, Peter [1996], "On the Direction of Technical Change," Discussion Paper A-506, University of Bonn.
- Funk, Peter [2005], "Entry and Growth in a Perfectly Competitive Vintage Model," *Journal of Economic Theory*, **138**: 211-236.
- Grossman, Gene and Elhanan Helpman [1991], "Quality Ladders in the Theory of Growth," *The Review of Economic Studies*, **58**: 43-61.
- Hannan, Timothy H. and John M. McDowell [1987], "Rival Precedence and the Dynamics of Technology Adoption: An Empirical Analysis," *Economica*, New Series, **54**: 155-171.
- Hellwig, M. and Andreas Irmen [2001], "Endogenous Technical Change in a Competitive Economy," *Journal of Economic Theory*, **101**: 1-39.
- Hirshleifer, Jack [1971], "The Private and Social Value of Information and the Reward to Inventive Activity," *The American Economic Review*, **61**: 561-574.
- Hopenhayn, Hugo [2006], "Knowledge, Diffusion, and Geography," plenary lecture, Society for Economic Dynamics.
- Irwin, Douglas A. and Peter J. Klenow [1994], "Learning-by-Doing Spillovers in the Semiconductor Industry," *The Journal of Political Economy*, **102**: 1200-1227.

- Jones, Larry and Rodolfo Manuelli [1990], "Convex Economic Growth," *The Journal of Political Economy* **98**, 1008-1038.
- Klevorick, A.K., R.C. Levin, R.R. Nelson, and S.G. Winter [1995], "On the Sources and Significance of Interindustry Differences in Technological Opportunities", *Research Policy* **25**, 185-205.
- Lerner, Joshua [2002], "Patent Protection and Innovation Over 150 Years", available at http://www.epip.eu/papers/20030424/epip/papers/cd/papers_speakers/Lerner_Paper_EPIP_210403.pdf, and from which we quote. Abridged version published as "150 Years of Patent Office Protection", *The American Economic Review Papers and Proceedings* **92**, 221-225.
- Makowski, Luis and Joseph M. Ostroy [1995], "Appropriation and Efficiency: A Revision of the First Theorem of Welfare Economics," *The American Economic Review* **85**, 808-827.
- Makowski, Luis and Joseph M. Ostroy [2001], "Perfect Competition and the Creativity of the Market," *Journal of Economic Literature* **XXXIX**, 479-535.
- Manuelli, Rody E. and Ananth Seshadri [2003], "Frictionless Technology Diffusion: The Case of Tractors," NBER Working Paper No. W9604.
- Park, Sangin [2002], "An Empirical Evaluation of the 1986 Semiconductor Trade Arrangement," mimeo SUNY Stony Brook.
- Plant, Arnold [1934], "The Economic Aspect of Copyright in Books," *Economica*, NS **1**(2), 167-195.
- Prescott, Edward and Ellen McGrattan [2008], "Openness, Technology Capital, and Development," paper presented at the University of Chicago Conference in Honor of Robert E. Lucas Jr.
- Quah, Danny [2002], "24/7 Competitive Innovation," mimeo, London School of Economics
- Rebelo, Sergio [1991] "Long Run Policy Analysis and Long Run Growth," *Journal of Political Economy*, **99**: 500-521.
- Romer, Paul M. [1990], "Endogenous Technological Change," *The Journal of Political Economy*, **98**: S78-S102.

- Rose, Nancy L. and Paul L. Joskow [1990], "The Diffusion of New Technologies: Evidence from the Electric Utility Industry," *RAND Journal of Economics*, **21**: 354-373.
- Sarkar, Jayati [1998], "Technological Diffusion: Alternative Theories and Historical Evidence," *Journal of Economic Surveys*, 12: 131 -176.
- Schumpeter, Joseph A. [1912], *Theorie der wirtschaftlichen Entwicklung* (transl. 1934, *The Theory of Economic Development: An inquiry into profits, capital, credit, interest and the business cycle*)
- Schumpeter, Joseph A. [1942], *Capitalism, Socialism and Democracy*, New York: Harper & Row
- Shapiro, C. [2001], "Navigating the Patent Thicket: Cross Licenses, Patent Pools and Standards Setting," in Jaffe, A., J. Lerner and S. Stern (eds) *Innovation and the Economy, Volume I*, MIT Press.
- Stigler, George J. [1956], "Industrial Organization and Economic Progress," in *The State of the Social Science* edited by Leonard D. White, University of Chicago Press, 269-282.
- von Neumann, John [1937], "Über ein ökonomisches Gleichungs-System und eine Verallgemeinerung des Brouwerschen Fixpunktsatzes," in K. Menger (ed.) *Ergebnisse eines Mathematischen Kolloquiums*, No. 8. Translated as "A Model of Economic General Equilibrium," *Review of Economic Studies* **XIII** (1945-1946).