Innovation and Competition with Non-Walrasian Labour Markets

Bryan Paterson
Department of Economics
Queen's University
paterson@econ.queensu.ca

Preliminary and Incomplete
Please do not cite

October 14, 2005

Abstract

This paper presents a model of innovation and growth in which firms choose positive levels of innovation activity in the presence of perfect product market competition both before and after a successful innovation. The reason why firms choose to innovate is not because of expected future profits in the product market, but because of expected future profits in the labour market. This result stems from the replacement of a standard Walrasian labour market with a market in which firms approach workers at random and bid for their services in an English auction format. A firm then chooses to engage in R&D in order to move to a higher productivity level so as to enhance its bargaining position relative to a potential worker and relative to rival firms bidding for the same worker's services. The expected surplus gained in the labour market by innovating is able to compensate the firm for its R&D expenditure.

1 Introduction

Models of endogenous growth through innovation typically imply that monopolies and other market structures of limited product market competition (PMC) are good for growth. The standard theoretical argument is that R & D activity is costly and so firms require some sort of expected future profit to compensate them for this cost. For example, if a firm that innovates successfully is granted a patent or is in some other way protected from PMC, then it can earn monopoly profits for a period of time, which then offsets what the firm had to incur in R & D costs in order to innovate in the first place. Since PMC reduces or removes monopoly rents altogether, then according to the theory it should be a hindrance to growth.

The problem is that most empirical work and 'real life' observations point to exactly the opposite conclusion given by the theory. For example, empirical papers like Nickell (1996) and Blundell, Griffith and Van Reenen (1999) find a positive link between PMC and innovation activity. In the last few years there have been several attempts to reconcile the theory of innovation with the empirical evidence, but most of these attempts involve detailed industry specific assumptions which would be difficult to apply to a general macroeconomic setting. One possible way to get generate innovation with PMC in the theory is by introducing the *escape competition effect*, whereby firms have a greater incentive to innovate in the face of greater competition because by moving to a higher technology level they can escape PMC with their rivals. The problem with this argument is that it still relies crucially on the fact that once a firm innovates, it must operate as a monopoly (or have a lot of market power) in order to recoup the costs incurred by its innovation activity.

This paper presents an alternative model of innovation and growth in which firms choose optimal (positive) levels of innovation activity even in the presence of perfect PMC both before and after a successful innovation. The reason why firms choose to innovate in this model is not because of expected future profits coming from the product market, but rather because of expected future profits coming from the labour market. In this paper the standard Walrasian labour market with a single market clearing wage is replaced by a more complex market where firms approach workers at random and bid for their services in an

English auction format. A firm then chooses to engage in R & D in order to move to a higher technology (productivity) level so as to enhance its bargaining position relative to a potential worker and relative to rival firms bidding for the same worker's services. The expected surplus a firm gains from the labour market by innovating (and moving to a higher technology level) is able to compensate the firm for its R & D expenditure. Hence, even perfect PMC as in this model is consistent with innovation and growth.

The outline of this paper is as follows. Section 2 gives a brief survey of some of the theoretical attempts to reconcile PMC with innovation. Section 3 presents a simple model of firms divided into two R & D coalitions. The firms engage in perfect competition in the product market, but they innovate collectively. The model yields a simple analytical solution that illustrates how expected returns in the labour market are sufficient to generate innovation and growth. In section 4 a more complex model is presented in which firms still engage in perfect PMC but instead engage in their own distinct R & D activities. In this case firms are distributed across technology levels and move individually from one technology level to another. Section 5 gives a brief discussion about the movement in wages over time implied by the model. Section 6 indicates the direction of future work on this paper.

2 Reconciling Innovation Activity with Product Market Competition

Aghion and Howitt (1998) suggest three possible theoretical strategies for breaking the negative link between PMC and innovation. The first of these strategies is to introduce agency considerations. By introducing individual agents within a firm, such as shareholders and managers, we can then find ways to relax the usual assumption that the goal of a firm is always to maximize profits. For example, while shareholders might want to maximize profits, a manager might want to maximize the probability that he retains his control of the firm (ie, keeps his job). Suppose that shareholders cannot easily observe how far away the firm is from its maximum profit level, but it is easy to observe how far away the firm is from the 'cutting edge' development of new products and services. In this case, a manager might maximize the probability of retaining his control of the firm by engaging in R & D in order to remain at the technological frontier. This may be his best strategy to convince the shareholders that he is doing a good job, but in general it will not be profit maximizing behaviour.

The second strategy proposed is to introduce a distinction between old and new product lines produced by firms, where the old product lines of one firm are (imperfect) substitutes for the new product lines of another firm. If we then define an increase in PMC as being an increase in the substitutability between old and new product lines, this could then encourage more R & D by firms to develop even newer product lines that are not as close substitutes with the old product lines of a rival firm.

The third strategy (and most promising of the three) is to replace *leap frog* innovation with *step by step* innovation. In a typical creative destruction model of endogenous growth and innovation, a firm that is a technological follower innovates in order to leap frog the leader and earn monopoly profits for a time until it is itself eclipsed by a future innovation from another firm. With step by step innovation, a technological follower must first catch up with the leader before it can pass the leader. This theoretical assumption allows two (or more) firms to be at the tech frontier at the same time. If we then define PMC as the degree of competition between these firms at the tech frontier, then PMC could encourage innovation activity through the *escape competition effect*; firms innovate in the presence of PMC in order to escape from competitive pressure from their rivals.

The problem with the first two strategies is that they are highly complex even at a micro level and any macro model that attempts to incorporate such details of firm organization and production would be intractable. The problem with the third strategy, as mentioned in the introduction, is that while it is able to generate innovation activity in the presence of PMC, this innovation activity depends critically on the expectation of future monopoly profits.

3 A Model of Two R & D Coalitions

Consider an economy with n firms and a large number of potential employees. The production technology of each firm is such that one worker is required to produce a final good which can then be sold in the product market. The firms in this economy are segmented into two groups. The first group consists of technological leaders who (after hiring a worker) can produce a final good with value $\overline{y}$. The second group of firms are technological followers, who can produce a final good with value $\underline{y}$, where $0 < \underline{y} < \overline{y}$.

Assume that there are at least two firms within each group and that consumers can identify the value (or quality) of a product offered by a firm before they buy it. Bertrand competition among firms in the product market ensures that firms at both technology levels will set prices for the goods they produce equal to their value; $\overline{p} = \overline{y}$ and $\underline{p} = \underline{y}$.

Suppose for a moment that every firm in this economy were operating at the low technology level. In that case firms might expend resources in order to innovate and become a technological leader, so long as the cost of their R & D would be offset by the expected future rents from being the first to innovate (and therefore being able to charge a monopoly price until another firm successfully innovates and becomes a rival at the higher technology level.)

This is the so called escape competition effect that can generate innovation activity in the presence of PMC. However, this effect will disappear the moment one firm is able to innovate and move from the low tech to the high tech level. No other firm at the low tech level will now have an incentive to innovate because they know that they will have to engage in Bertrand competition at the high tech level and as a result there will be no rents to reward continued R & D expenditure. Innovation in the economy as a whole drops to zero once one firm innovates.

However, we can sustain innovation activity in this economy even with perfect PMC at both the high and low technology levels, if there are rents that can be extracted from the labour market by moving from the low tech level to the high tech level. For example, suppose that firms in this economy can each approach one potential employee in the labour market (at random.) Each worker could then be contacted by zero firms, by one firm or by more than one firm. If a worker is not contacted by a firm, he gets a payoff of 0 (which is assumed to be his outside option.) If a worker is contacted by one or more firms, then the wage w he gets is determined by an English auction, where $w \geq 0$.

Now index the wage of a particular worker according to $w(i, j)$, where $i \in \{0, \underline{y}, \overline{y}\}$ is the highest tech level in the set of firms that have contacted him and $j \in \{0, \underline{y}, \overline{y}\}$ is the second highest tech level among the firms that have contacted him. If $i = 0$ then a worker has not been contacted by any firms and if $j = 0$ then a worker has only been contacted by one firm. With wages determined by an English auction, the distribution of wages in the economy will be as follows:

$$w(0,0) = 0$$

$$w(\underline{y}, 0) = 0$$

$$w(\underline{y}, \underline{y}) = \underline{y}$$

$$w(\overline{y}, 0) = 0$$

$$w(\overline{y}, \underline{y}) = \underline{y}$$

$$w(\overline{y}, \overline{y}) = \overline{y}$$

Notice first that if a worker is approached by only one firm, then the firm has all of the bargaining power and retains all of the surplus from production (so that the worker gets 0.) Second, the advantage of being a high tech firm is seen especially in $w(\overline{y}, \underline{y}) = \underline{y}$. A high tech firm that approaches a worker with offers from other low tech firms can still retain some of the surplus from production by offering the maximum wage that the low tech firms can afford to pay.

Using this wage distribution, we can now determine the expected profits of firms at both the low tech and high tech levels. Let $\hat{\alpha}(n)$ denote the probability that a firm contacts a worker that does not have an offer from any other firm. Let $\alpha(n, \theta)$ denote the probability that a high tech firm contacts a worker that only has one (or more) offers from a low tech firm (and no offers from other high tech firms.) Here θ is the fraction of the total number of firms that are at the high tech level. We also assume that there are no lasting employment relationships between firms and workers; a worker is hired, produces for the firm a good of value $\underline{y}$ or $\overline{y}$ depending on the technology level of the firm, earns a wage and is then fired and returns to the labour market. The expected profit flows of low and high tech firms will then be given by

$$\pi(\underline{y}; n) = \hat{\alpha}(n)\underline{y} \tag{1}$$

$$\pi(\overline{y}; n, \theta) = \hat{\alpha}(n)\overline{y} + \alpha(n, \theta)(\overline{y} - \underline{y}) \tag{2}$$

Assuming that a new firm must pay a fixed cost k to enter this economy and that entry can only occur at the low tech level, n will be determined by the free entry condition $\hat{\alpha}(n^*)\underline{y} - k = 0$.

The next step is to model the innovation activity of the low tech and high tech firms. First, assume that there is an equal number of firms at both tech levels (ie, $\theta = \frac{1}{2}$) and that the two groups of firms form a coalition to engage in R & D as a whole. For example, while all the low tech firms compete with each other in the product market, they cooperate in their innovation activity by choosing a joint level of R & D expenditure that is shared equally among them. If the innovation is successful, then all low tech firms advance to the high tech level.

Second, assume that if the high tech firms innovate successfully and move to an even higher technology level given by $y' > \overline{y}$, then the low tech firms behind them are able to immediately adopt the 'old' technology of the high tech firms, which is $\overline{y}$. In other words, the maximum sustainable gap in technology between the two groups of firms is one step. This ensures that so long as the gap between y' and $\overline{y}$ and between $\overline{y}$ and $\underline{y}$ is the same, high tech firms will have no incentive to innovate unless the low tech firms have innovated successfully and moved to the high tech level (so that all firms are at the same tech level.)

Finally, assume that the cost to a coalition of firms of choosing an R & D intensity equal to a is given by $\frac{a^2}{2}$, where $a \in [0, 1]$. Now we can write the value functions for the high and low tech coalitions of firms as follows:

$$rV_1 = \frac{n}{2}(\hat{\alpha}\overline{y} + \alpha(\overline{y} - \underline{y})) + a_{-1}(V_1 - V_0) \tag{3}$$

$$rV_{-1} = \frac{n}{2}\hat{\alpha}\underline{y} + a_{-1}(V_0 - V_{-1}) - \frac{a_{-1}^2}{2} \tag{4}$$

$$rV_0 = \frac{n}{2}\hat{\alpha}\underline{y} + a_0(V_1 - V_0) + \bar{a_0}(V_{-1} - V_0) - \frac{a_0^2}{2} \tag{5}$$

From equation (3), we see for example that the flow value of a coalition that is at the high tech level is equal to the sum of the current profit flows of each firm in the coalition plus the expected capital loss that will result when the low tech coalition innovates successfully and moves to the high tech level (so that both coalitions are at the same tech level.) In a similar way, equation (4) gives the flow value of a coalition that is at the low tech level and equation (5) gives the flow value to both coalitions when they are at the same tech level. The first order conditions with respect to a_0^* and a_{-1}^* are then given by

$$a_0^* = V_1 - V_0$$

$$a_{-1}^* = V_0 - V_{-1}$$

Substituting these expressions back into the value functions above and simplifying we can derive the following optimal values for a_0^* and a_{-1}^*:

$$a_0^* = -r + \sqrt{r^2 + n(\hat{\alpha} + \alpha)(\overline{y} - \underline{y})} \tag{6}$$

$$a_{-1}^* = -(r + a_0) + \sqrt{r^2 + (a_0)^2 + n(\hat{\alpha} + \alpha)(\overline{y} - \underline{y})} \tag{7}$$

First notice that innovation expenditure when all firms are at the same technology level, given by a_0^*, is increasing in $(\hat{\alpha} + \alpha)(\overline{y} - \underline{y})$, which is the expected gain in the labour market to being the technological leader. A coalition of firms will engage in costly R & D because a successful innovation will increase their expected surplus in the labour market, even though they will be perfect competitors with each other in the product market. The greater the gap between $\overline{y}$ and $\underline{y}$, the greater the reward to innovation and so the higher the R & D expenditure when the two coalitions of firms are level with each other.

Innovation by the technological followers a_{-1}^* does not occur because of any direct benefit from catching up with the technological leaders, since the current profit flows in (4) and (5) are the same. Rather a_{-1}^* is positive because a coalition of firms lagging behind its rival must catch up with its rival before it can be in a position to innovate again and enjoy the expected surplus in the labour market that results from being a technological leader. This result stems from the assumption that technological innovation is of the step by step form rather than the leap frog form used in more common creative destruction models.

4 A Model of Multiple Firms Engaged in R & D

The analysis in the previous section is simplified greatly by assuming that all firms can be divided into two rival R & D coalitions. This allows perfect competition to be maintained in the product market, but at the same time it reduces the optimal R & D analysis to a duopoly problem. This in turn means that optimal innovation expenditure depends only on a coalition's technological position relative to its rival.

Once we move away from this simple framework, optimal innovation by a firm will depend not only on its technological position relative to other firms, but also on the distribution of firms across the technology levels. As before we assume that there can never be more than two technology levels at any one time. With this assumption we can then describe the distribution of firms across technology levels by θ, which will now in general be different from $\frac{1}{2}$.

If a firm that is currently a technological leader at $\overline{y}$ innovates successfully and moves to a higher tech level, y', then all other firms that were at the old tech frontier $\overline{y}$ become tech followers and all firms that were tech followers at $\underline{y}$ are assumed to be able to immediately catch up to $\overline{y}$ (so that no firm is ever more than one step behind the technological leader(s).) This process is equivalent to all firms at $\overline{y}$ moving down to $\underline{y}$ except for the tech leader firm that innovates successfully.

Because of this assumption, firms that are tech followers and firms that are tech leaders will never innovate at the same time. If firms that are tech leaders are engaging in R & D, then firms that are tech followers observe this behaviour and set their innovation expenditure to zero (since once a single tech leader firm is successful, all tech followers get a free ride up to a higher tech level.)

The dynamics of the technological advance of firms is as follows. Suppose initially that θ is small so that there is a small fraction of firms at $\overline{y}$ with most firms operating as tech followers at $\underline{y}$. The firms that are tech leaders enjoy high expected surplus in the labour market given by $\alpha(\theta))(\overline{y}-\underline{y})$ where α is now a decreasing function of θ. Firms that are tech followers engage in their own individual R & D activities in order to innovate and move successfully from $\underline{y}$ to $\overline{y}$, so that they can then enjoy added expected surplus in the labour market. However, this extra expected surplus for the tech leader firms decreases as more and more firms innovate successfully and move from $\underline{y}$ to $\overline{y}$ (because α falls as θ increases.) Define $\hat{\theta}$ as the critical fraction of firms at the tech frontier such that it becomes optimal for tech leader firms to attempt to innovate again and move to a higher tech level. At this point, all remaining tech follower firms will set their R & D expenditures to zero and once a tech leader firm has innovated successfully, the entire process just described begins again.

Suppose then that $\theta < \hat{\theta}$ so that the technological followers are innovating (and the technological leaders are not.) We can then write the value function for a tech leader firm as

$$rV_1(\theta) = \hat{\alpha}\overline{y} + \alpha(\theta)(\overline{y} - \underline{y}) + q(\theta)(V_1(\theta') - V_1(\theta)) \tag{8}$$

where $q(\theta)(V_1(\theta') - V_1(\theta))$ is the expected capital loss for a tech leader firm when tech follower firms innovate successfully and move from $\underline{y}$ to $\overline{y}$. This will increase θ to θ', which will in turn decrease the firm's expected surplus in the labour market.

The value function for a tech follower firm is given by

$$rV_{-1}(\theta) = \hat{\alpha}\underline{y} + a_{-1}(V_1(\theta') - V_{-1}(\theta)) + \bar{q}(\theta)(V_{-1}(\theta') - V_{-1}(\theta)) - \beta\frac{a_{-1}^2}{2} \tag{9}$$

where $a_{-1}(V_1(\theta') - V_{-1}(\theta))$ is the expected capital gain from moving from a tech follower to a tech leader, and $\bar{q}(\theta)(V_{-1}(\theta') - V_{-1}(\theta))$ is the expected capital loss from being left behind in the innovation race as other tech followers become tech leaders. β is a parameter that can be set to ensure that $0 \leq a_{-1} \leq 1$.

The first order condition associated with R & D expenditure for tech followers is given by

$$a_{-1}^*(\theta) = \frac{1}{\beta}(V_1(\theta') - V_{-1}(\theta))$$

Because nearly everything in this case depends on the distribution of firms across the two technology levels, it is not possible to derive an analytical solution for $a_{-1}^*(\theta)$ as in the previous section. However by substituting the first order condition back into the two value functions above, we can derive an intuitive expression for $a_{-1}^*(\theta)$. If we let $\beta = 1$ then we can arrive at the following:

$$a_{-1}^* = -r + \sqrt{r^2 + 2[(\hat{\alpha} + \alpha(\theta))(\overline{y} - \underline{y}) + q(\theta')(V_1(\theta'') - V_1(\theta')) - \bar{q}(\theta)(V_{-1}(\theta') - V_{-1}(\theta))]}$$

Note first of all that a_{-1}^* is increasing in expected profit flow from being at the tech frontier, given by $\hat{\alpha} + \alpha(\theta))(\overline{y} - \underline{y})$. Second, we see that a_{-1}^* is decreasing in the expected capital loss that will occur once a firm moves to the tech frontier and then more low tech firms move to the same tech frontier in the future, given by $q(\theta')(V_1(\theta'') - V_1(\theta'))$. Third, a^*_{-1} is increasing in the expected capital loss that will occur as other firms move to the higher tech level before you, given by $\bar{q}(\theta)(V_{-1}(\theta') - V_{-1}(\theta))$

Finally, it is important to note that a^*_{-1} is decreasing in θ. That is, as more and more tech followers become tech leaders, there is less incentive for the remaining tech followers to innovate. This property suggests that overall innovation activity (which is a proxy for growth) moves in *cycles*. Innovation activity is very high when θ is low, since there are many tech followers innovating at high levels. Over time, as more and more firms innovate successfully and move from $\underline{y}$ to $\overline{y}$, overall innovation activity falls (since there are fewer firms innovating and those that are innovating are doing so at lower intensities.) This slowing of innovation activity continues until $\theta = \hat{\theta}$, at which point overall R & D jumps to a new higher level as all firms at the tech frontier begin to attempt to innovate.

Suppose then that $\theta > \hat{\theta}$ so that the technological leaders are innovating. Then the value function for a tech leader is given by

$$rV_1(\theta) = \hat{\alpha}\overline{y} + \alpha(\theta)(\overline{y} - \underline{y}) + a_1(V_1(\theta') - V_1(\theta)) + \bar{s}(\theta)(V_{-1}(\theta') - V_1(\theta)) - \beta\frac{a_1^2}{2} \tag{10}$$

where $a_1(V_1(\theta') - V_1(\theta))$ is the expected capital gain from moving from the existing tech frontier to a new tech frontier (where with fewer firms, $\theta' < \theta$), and $\bar{s}(\theta)(V_{-1}(\theta') - V_1(\theta))$ is the expected capital loss from losing the innovation race to a rival firm and slipping from the tech frontier and becoming a tech follower.

The value function for a tech follower in this case will be

$$rV_{-1}(\theta) = \hat{\alpha}\underline{y} + s(\theta)(V_{-1}(\theta') - V_{-1}(\theta)) \tag{11}$$

where $s(\theta)(V_{-1}(\theta') - V_{-1}(\theta))$ is the expected capital gain from the free ride to a new tech follower position once a tech leader innovates successfully.

The first order condition associated with R & D expenditure for tech leaders is given by

$$a_1^*(\theta) = \frac{1}{\beta}(V_1(\theta') - V_1(\theta))$$

If we let $\beta = 1$ and we substitute this first order condition into the two value functions above, we can arrive at the following expression for the R & D expenditure for tech leaders

$$a_1^* = -r + \sqrt{r^2 + 2[(\alpha(\theta') - \alpha(\theta))(\overline{y} - \underline{y}) + \Delta\bar{s}(\theta)(V_{-1}(\theta') - V_1(\theta))]} \qquad (12)$$

Note first that a_1^* is increasing in the expected profit flow from being at a new tech frontier (with fewer firms at this tech frontier), given by $(\alpha(\theta') - \alpha(\theta))(\overline{y} - \underline{y})$.

Second, note that a_1^* is decreasing in the change in expected capital loss that will result when a firm slips from the high tech level to the low tech level, which is given by $\Delta\bar{s}(\theta)(V_{-1}(\theta') - V_1(\theta))$

Finally, it can be shown (and is intuitive) that a_1^* is increasing in θ. We can now define the critical value for $\theta = \hat{\theta}$ more precisely, where $\hat{\theta}$ is such that $a_1^* = 0$ or

$$\alpha(\theta') - \alpha(\theta^*))(\overline{y} - \underline{y}) = \Delta\hat{q}(\theta^*)(V_{-1}(\theta') - V_1(\theta^*))$$

In other words for $\theta \leq \hat{\theta}$, $a_1^* = 0$ and for $\theta > \hat{\theta}$, $a_1^* > 0$.

We know that if $a_1^* > 0$ then $a_{-1}^* = 0$, since as soon as one of the high tech firms innovates and moves to a higher tech level, all firms will move immediately to the tech level below the new tech frontier. As mentioned earlier, firms that are currently at the low tech level will have no incentive to innovate because they will get a free ride up to a higher technology level.

One interesting case that could arise is if there are values for θ for which both $a_1^* = 0$ and $a_{-1}^* = 0$. In this case, the distribution of firms across the two technology levels would be such that there would be too many firms at the tech frontier to give an incentive for tech followers to innovate, and at the same time there would be too few firms at the tech frontier to give an incentive for the existing tech leaders to innovate. In such a case, innovation (and growth) would grind to a halt. Without introducing a fixed cost to innovation, it is unlikely that both R & D intensities would ever both be zero. However, it could be that for a range of $\theta < \hat{\theta}$, a_{-1}^* could fall to very low levels, so that innovation activity (and growth) could stagnate for an extended period of time.

5 The Evolution of Wages Over Time

As discussed above, this model predicts that total innovation activity (as a proxy for economic growth) in the economy will move in waves. Average productivity (defined as the average tech level of firms) will rise over time. What does the model imply about the movement of wages over time?

Recall from section 2 that workers in the labour market can be offered (and earn) an instantaneous wage $w \in \{0, \underline{y}, \overline{y}\}$. The fraction of workers in the labour force that earn a wage of 0 depends on the relative number of firms and workers in the labour market. Recall that the number of firms in the market is given by the free entry condition $\hat{\alpha}(n^*)\underline{y} - k = 0$, where k is the fixed cost to entry. As technological advance occurs over time, $\underline{y}$ will rise and the number of firms will increase. If the number of firms increases faster than the number of potential workers, this will decrease the probability that a worker is approached by only one (or zero) firms, thereby reducing the frequency in which a worker earns a wage of 0.

When θ is low (so that there are few firms at the tech frontier), very few workers will earn $w = \overline{y}$, since this requires that a worker be approached by two firms at the tech frontier. Most workers will earn $w = \underline{y}$ or 0. The average wage will rise over time as more firms innovate successfully and become tech leaders, increasing the probability that a worker is approached by two firms at the tech frontier. Once the tech leaders begin to innovate and one (or more) of them is successful, there is an immediate jump in the technology of all tech follower firms to the old tech frontier. This in turn will generate a distinct jump in the average wage.

In general wages will lag behind technology (productivity) improvements, but will follow the same general dynamics.

6 Further Work

The first task will be to solve the more complex model of section 4 numerically in order to better understand the innovation behaviour of firms and the movement of firms across technology levels. This in turn will generate more precise descriptions of how wages and the wage distribution changes over time.

The second (and more ambitious) task is to relax the strong assumption that there can only be two technology levels at any one time. Thought must be given to how to model a group of firms that may disperse themselves across a number of different technology levels. Intuitively, this will smooth out the distinct 'jumps' in the average technology level and in the average wage that occur when there is an automatic catch up by the tech followers after a tech leader innovates.

References

[1] Aghion, P., N. Bloom, R. Blundell, R. Griffith and P. Howitt (2002), *Competition and Innovation: An Inverted U Relationship*, NBER Working Paper 9269.

[2] Aghion, P., C. Harris, P. Howitt and J. Vickers (2001), *Competition, Imitation and Growth with Step-by-Step Innovation*, **Review of Economic Studies**, vol.68, pp.467-492.

[3] Aghion, P. and P. Howitt (1998), **Endogenous Growth Theory**. Cambridge, MA: MIT Press.

[4] Aghion, P., C. Harris and J. Vickers (1997), *Competition and Growth with Step-by-Step Innovation: An Example*, **European Economic Review, Papers and Proceedings**, pp.771-782.

[5] Blundell, R., R. Griffith and J. Van Reenen (1999), *Market Share, Market Value and Innovation in a Panel of British Manufacturing Firms*, **Review of Economic Studies**, vol.66, pp.529-554.

[6] Julien, B., J. Kennes and I. King (2001), *'Residual' Wage Disparity in Directed Search Equilibrium*, manuscript, University of Auckland.

[7] Nickell, S. (1996), *Competition and Corporate Performance*, **Journal of Political Economy**, vol.104, pp.724-746.

[8] Postel-Vinay, F. and J.-M. Robin, (2002), *The Distribution of Earnings in an Equilibrium Search Model with State Dependent Offers and Counteroffers*, **International Economic Review**, vol.43, pp.989-1016.