


Innovation, Imitation and Open Source

Rufus Pollock

EERI Research Paper Series No 20/2008



EERI

Economics and Econometrics Research Institute

Avenue de Beaulieu

1160 Brussels

Belgium

Tel: +322 299 3523

Fax: +322 299 3523

www.eeri.eu

INNOVATION, IMITATION AND OPEN SOURCE

RUFUS POLLOCK

UNIVERSITY OF CAMBRIDGE

AUGUST 2008

ABSTRACT. An extensive empirical literature indicates that, even without formal intellectual property rights, innovators enjoy a variety of first-mover advantages and that ‘imitation’ is itself a costly activity. There is also accumulating evidence that an ‘open’ approach to knowledge production can deliver substantial efficiency advantages. This paper introduces a formal framework incorporating all of these factors. We examine the relative performance of an ‘open’ versus a ‘closed’ (proprietary) regime, and explicitly characterise the circumstances in which an open approach, despite its effect on facilitating imitation, results in a higher level of innovation.

Keywords: Innovation, Imitation, Intellectual Property, Openness, Open Source

JEL codes: L17, L5, O3

1. INTRODUCTION

The last decade or so have seen an explosion in the level of ‘open’¹ information production not only in traditional areas such as software (e.g. Linux, Firefox) but also in areas ranging from online encyclopaedias (Wikipedia) to genomic databases (the Human Genome Project). Such developments suggest that, at least in some important cases, open models of knowledge production can do as, or even better, than closed² ones. As we discuss

Faculty of Economics and Emmanuel College, University of Cambridge. Email: rp240@cam.ac.uk or rufus@rufuspollock.org. Licensed under a Creative Commons Attribution license (unported) v3.0. I thank Rupert Gatti and David Newbery for valuable input on early versions of these ideas. I also wish to thank participants at the 2008 FLOSS 2.0 workshop for their comments and suggestions.

¹For the purposes of this paper a work is ‘open’ or ‘open-source’ if it can be used, redistributed and reused freely. Here, freely means without monetary or other substantial restriction though, for example the requirement of attribution or even that derivative works be re-shared, does not render a work unfree. However it does exclude the requirement of payment or that explicit permission be sought, or the imposition of restrictions on the type of use (such as limiting it to research or non-commercial activities). Furthermore, since, without access, a piece of knowledge cannot be used it also excludes the use of secrecy – ‘open’ knowledge must be publicly available.

²That is those based on on secrecy, exclusive rights (IP) or some combination thereof.

further below, this is somewhat surprising and it is important to investigate carefully how and why this could be the case.

It is important to note here that we are focused on the rate of innovation and not the level of welfare. After all it is well-accepted that being more ‘open’ (having weaker intellectual property rights) can improve welfare by improving access. But this is certainly not the case in relation to innovation. In fact most of the traditional literature on innovation would support, implicitly or explicitly, reasoning along the following lines: the level of innovation is a function of the value that innovators (that is the original creators of a work or invention) obtain, V , and the cost they incur, F . Furthermore, it is clear that the level (or likelihood) of innovation is increasing in the value and decreasing in cost. Label the closed and open regimes by C and O respectively then it is usual to assume that the proprietary regime results in higher returns than in the open one – $V^P > V^O$ – while costs are unchanged – $F^P = F^O$. This necessarily implies that innovation must be higher in the proprietary regime than in the open one. Furthermore, suppose imitation has cost F^{IMM} but is ‘fast’.³ Then, simple competitive free entry arguments would suggest that the common returns (common because imitation is fast) must be driven down to costs so the innovator receives $V = F^{IMM}$. Since an innovator must cover her costs innovation can only occur if imitation is at least as costly as innovation: $F^{IMM} \geq F$. But, at least in the open case, imitation is certain to be cheaper than innovation and hence no innovation will occur!

This brief sketch of the standard approach already indicates why the level of ‘open’ production is rather surprising. It also points us to where we must look if we are to explain the success, or simply survival, of open approaches: either we can examine costs (F) or income/value V . At this point it is worth recalling that repeated surveys, such as Levin et al. (1987), Mansfield (1985), Cohen et al. (2000), and Arundel (2001), show that firms appropriate returns from innovation using a variety of methods other than exclusive rights (IP) such as secrecy, lead time, marketing and sales, learning curve advantages. Not

³It is noteworthy that much of the existing theoretical literature has tended to assume ‘perfect’ nonrivalry, that is, that an innovation (or creative work) once made may be costlessly, and instantaneously, reproduced if not protected by IP or hidden from others via secrecy. For example, Nordhaus (1969) (and following him Scherer (1972)), in what is considered to be one of the founding papers of the policy literature, implicitly assume that without a patent an innovator gains no remuneration. Similarly, Klemperer (1990) in his paper on patent breadth makes clear his assumption of costless imitation as do Scotchmer and Green (1990). For a more recent example see e.g. Hopenhayn and Mitchell (2001).

only does this indicate that there are a variety of first-mover advantage for an innovator, but this work also shows that imitation is a costly process both in terms of time and money even when no exclusive rights are used (or available). Of course, the major alternative to IP is often secrecy. In the case of ‘open’ (‘open source’) knowledge production both secrecy and the traditional exclusivity of IP are foregone. Nevertheless, even in this case, it is likely that still exist some degree of imitation costs and first-mover advantage for the innovator.

This paper takes this empirical evidence ‘seriously’ and we explicitly model imitation as costly (though though still usually as cheaper than innovation) and the innovator as having some form of first-mover advantage. Our framework provides a simple, and novel, way to conceptualize ‘innovation’ space, which allows one to compare innovation under a proprietary and an ‘open-source’ regime in a straightforward and intuitive manner. We show that, even in the baseline case where ‘open-sourcing’ simply reduces imitation (copying) costs, some innovation will still take place under the open-source regime – albeit substantially less than in the proprietary case. The crucial point here is that at low innovation costs ‘allowable’ imitation costs (that is imitation costs that still result in innovation being made) can be very low and thus even for large reductions in the cost of imitations some innovations remain feasible.

Our next step is to consider the possibility that an ‘open’ approach reduces both innovation and imitation costs. A variety of authors (e.g. Lakhani and von Hippel (2003); Bessen (2006)) have pointed out that an ‘open-source’ approach may offer substantial efficiency advantages – for example by allowing users to participate directly in adding features and fixing ‘bugs’ – and this is particularly true where the information good is complex and (hence) transaction costs are high. While the initial arguments in favour of the efficiency of open approaches were based on anecdotes or case study recent empirical work on a larger scale, such as Koch (2008), have provided strong empirical support for his view. Thus, we explicitly model the move to ‘open-source’ as resulting in a simultaneous reduction in *both* innovation and imitation cost. This clearly makes an open approach more attractive compared to the baseline case. However, the crucial question is: to what extent – after all the situation may well still be (far) inferior to that under a closed approach. Here, we are able to situate the effect of a change of regime within a single overarching, but simple,

framework. Using it, we can compare, clearly, and formally, different assumptions about what going ‘open’ entails, and derive specific (sufficient) conditions under which a move to an open regime results in an increase in innovation.

2. THE MODEL

There is some information good, the ‘innovation’,⁴ which can be created at fixed cost F_i – the innovation cost. Once created, it may be imitated with cost F_m (implicitly assumed to be common across all imitators).⁵

What does this imitation cost, F_m , represent? The empirical literature discussed above suggests two main options: either it represents the cost of imitation in money or the cost in time.⁶ In general we need not be too concerned with whether it is one of these, or the other – or some combination of both. Rather what will matter for our purposes is that the imitation ‘cost’ a) affects the income of the innovator b) is itself affected by the regime (open or proprietary) under which the information good is produced.

Specifically, define Π to be gross profits of the innovator (that is revenues net of all expenses other than the main innovation cost). Then revenues are a function of imitation cost: $\Pi = \Pi(F_m)$, furthermore as imitation costs rise innovator’s profits are non-decreasing: $\Pi' \geq 0$. Finally, though infrequently used in what follows, it will be useful to define $\Pi_m(F_m)$ as the profit of the imitators.⁷

In what follows, we will rarely go beyond this reduced form approach to specify the exact mechanism by which imitation costs impact on the innovator’s income as this allows us to proceed with greater generality. For our purposes, all we will require is that costs and profits are all denominated in comparable terms – for example, in terms of a numeraire good ‘money’. Nevertheless, it may be useful for the reader to have in mind some concrete examples of the underlying strategic interactions by which an innovator’s profits are determined.

⁴Throughout we will use the term innovation and imitation rather than alternatives such as creator and copier. Though the former are usually more associated with (patentable) ‘ideas’ while the second are more associated with (copyrightable) expressions no such distinction is intended here.

⁵This model can easily be generalized to allow variations in imitation costs – for example imitation costs could fall as more and more imitators enter, or it could fall with time etc.

⁶There are clearly other possibilities, for example imitation may be limited by the availability of skilled labour, or access to other necessary complementary assets – see e.g. Teece (1986). However, these are both more complex to model and, we believe, of lesser importance than the main factors of time and money.

⁷If profits vary across imitators define this as the maximum of the profits earned by imitators.

One possibility is that F_m represents pure monetary cost and the innovator and imitators compete to supply a homogenous good in Stackelberg fashion (the innovator being the first-mover).⁸ In that case, since imitators must cover their own fixed cost, F_m defines, via a zero profit free-entry condition, the total number of imitators who enter and hence the price level. Lower imitation costs lead to more entry and hence lower prices.

Another possibility is that F_m represents time, that is, the delay before imitators are able to produce a competing product. In this case the innovator enjoys a monopoly prior to imitator entry followed by some subsequent competitive game. Again it is clear that a lower imitation ‘cost’ corresponds to lower profits for an innovator. Yet another possibility is to have competition between innovators and imitators but with differentiated products (whether vertically or horizontally differentiated). Here, F_m would be some measure of the proximity of imitators to the innovator in product space and as F_m rises innovator’s enjoy a larger/stronger local monopoly and hence once again have larger (gross) profits.

2.1. A Normalization. Let M be monopoly profits for the innovator (i.e. without any imitative entry). No agent’s profits can be greater than monopoly profits M . Hence, it makes sense to simplify by normalizing profits and innovation fixed costs by dividing them by M :

$$\begin{aligned} f_i &= \frac{F_i}{M} \\ f_m &= \frac{F_m}{M} \\ \pi &= \frac{\Pi}{M} \end{aligned}$$

We can also take (normalized) profits π to be a function of f_m rather than F_m . It is also useful to define the proportional imitation cost $\phi = \frac{F_m}{F_i} = \frac{f_m}{f_i}$.

2.2. The Space of Innovations. In this model an innovation (under a particular regime – see next section) is specified by the tuple consisting of its ‘innovation’ cost and its ‘imitation’ cost: (f_i, f_m) . This conveniently allows us to visualize innovation space in a two dimensional graph (see the figures below for examples). Specifically, innovation and imitation costs are non-negative, $f_i, f_m > 0$ and under our normalizations must lie between

⁸Pollock (2006) considers this particular model in detail.

0 and 1. Thus, under the assumptions given and using normalized variables the space of innovations is $IS = \{(f_i, f_m) \in [0, 1] \times [0, 1]\}$, that is the unit square.

In addition, we note that if imitation cost is less or equal to innovation cost, as is often assumed, we have: $f_m \leq f_i$. In this case no innovation will ever lie above the ‘diagonal’ in the unit square and $IS = \{(f_i, f_m) \in [0, 1] \times [0, 1] : f_m \leq f_i\}$. Under the intuitive equation of area with amount, using (f_i, f_m) coordinates gives the impression that the number of (potential) innovations is dropping as innovation cost drops (since the height under the diagonal is falling with innovation costs). This is rather misleading and to correct this it is natural to change to (f_i, ϕ) coordinates in which case innovation space is again the full unit square: $IS = \{(f_i, \phi) \in [0, 1] \times [0, 1]\}$. Obviously there is a one-to-one mapping between these coordinate systems and so it does not strictly matter which we use. However, in what follows it will frequently be more convenient to work with (f_i, ϕ) rather than (f_i, f_m) .

2.3. Proprietary and Open Regimes. We shall be interested in comparing different regimes under which a given innovation (or set of innovations) is produced. A regime (R) can have two distinct effects. First, it can affect the innovation and imitation costs (f_i, f_m) , that is, the location of an innovation in ‘innovation space’ (IS). In particular, location of an innovation (or set of innovations) can be represented by some density function on IS, say g which is a function of the regime – so $g = g^R$.⁹ Second, the regime may alter the profit function itself π , that is, the manner in which imitation cost affects the (gross) profits of the innovators.

Here, we focus on the first of these effects only and assume that the regime has no impact on the profit function itself – of course, it will affect profits since these are a function of imitation costs. We will be interested in comparing and contrasting two particular regimes: a closed/proprietary one – based on some combination of secrecy and exclusive rights (e.g. patent or copyright) – and an open one. These two regimes, via their impact on imitation and innovation cost will lead to the same innovation (or set of innovations) being located in different positions in innovation space. This is illustrated in Figure 1 for the case where the move from closed to open results in a drop in both innovation costs f_i and a reduction

⁹This distributional approach gives us a natural way to represent not only a single innovation or technology but a whole set of innovations – for example at the industry or economy wide level – as well to represent uncertainty with respect to the location of a particular innovation or technology.




FIGURE 1. Effect of different regimes on the location of an innovation or innovation set. As shown, an open regime results in both lower innovation and lower imitation costs (both absolutely, f_m , and proportionately ϕ).

in proportional imitation costs ϕ (which corresponds to an even larger drop in imitation costs since $f_m = \phi f_i$).

3. ANALYSIS

Our analysis proceeds in stages. First, we define the feasible set, that is the subset of innovations in innovation space that would occur with a particular profit function. Next, we deduce some basic properties and structure of this feasible set. Given a particular set of (potential) innovations a regime determines their location in innovation space. Intersecting that location with the feasible set we have the set of innovations that occur under that particular regime. Since, as just discussed above, the profit function does not itself change across regimes the feasible set remains constant. Thus, finally, we look out how a change in regime affects the level of innovation by examining how the regime alters the location of (potential) innovations within innovation space, and thereby alters the intersection of these innovations with the feasible set. In particular, we can provide non-trivial conditions on the size of the reductions in innovation and imitation cost under which a move to an open regime increases the level of innovation.

3.1. The Feasible Set. Take a particular regime R . This has associated with it some (gross) profits function: $\pi(f_m)$. A (potential) innovation, (f_i, f_m) , is feasible if, and only if, it generates a non-negative net surplus for its associated innovator: $\pi(f_m) - f_i \geq 0$ – that is if, and only if, its fixed cost is less than its (gross) profits. We may then define the ‘feasible set’ as the subset of IS consisting of all (potential) innovations which are feasible:

Definition 1 (The Feasible Set). The ‘Feasible Set’, F (or F_R , for a regime R is the subset of innovation space whose corresponding innovations occur under R :

$$F = \text{Feasible Set} = \{(f_i, f_m) \in IS : f_i \leq \pi(f_m)\} = \{(f_i, \phi) \in IS : f_i \leq \pi(\phi f_i)\}$$

What can we say about the structure of the feasible set? First, observe that the feasible set is closed and hence it includes its boundary $\bar{F} = \partial F$. Next note that, with innovation cost on the x-axis in the IS diagram, the feasible set will simply be the space lying to the left of the line $f_i = \pi(f_m)$. Since f_i is plotted on the x-axis it is natural to transform this to a function of innovation cost: F is then the set lying above the line $g(f_i)$ where $g = \pi^{-1}$. In terms of proportional imitation cost coordinates (f_i, ϕ) this is the line given by $h(f_i)$ where $h = g(f_i)/f_i$, and the feasible set are those points lying above this line.¹⁰

Thus, what is needed is to characterise how gross profits, $\pi(f_m)$, vary with imitation costs. First recall that we have $\pi' > 0$ that is gross profits are increasing in imitation cost. Next, observe that as we are using renormalized variables $\lim_{f_m \rightarrow 0} \pi(f_m) = 1$ and so $(f_i, 1) \in F \forall f_i \in (0, 1)$ (in particular, $(1, 1)$ is in F). Similarly, the limit of profits as imitation costs go to zero must be non-negative: $\pi(0) \geq 0$. Figure 2 shows a graphical rendering of one possible feasible set conforming to these restrictions together with a labelling of the main features (e.g. g, h curves). To go beyond these basic observations we need to introduce a condition on the profit function. First a definition:

Definition 2. A function $H(x)$ is super-linear (sub-linear) if $H(kx) \geq kH(x)$ ($H(kx) \leq kH(x)$) $\forall k \in (0, 1)$. That is H lies below (above) any ray from the origin to x .

Note that a sufficient condition for H to be super-linear (sub-linear) is that (a) $H(0)$ is non-negative (b) H is quasi-concave (quasi-convex) (Pf: take $0, x$ as the endpoints of

¹⁰Note that while g was increasing – as innovations cost rise so the feasible imitation cost must rise – h may not be.




FIGURE 2. An example of a feasible set corresponding to $\pi(f_m) = \sqrt{f_m} \Rightarrow g(f_i) = f_i^2, h(f_i) = f_i$. The functions shown conform to the restrictions described in the text (in fact they derive from the Stackelberg example already discussed in the section 2).

an interval and apply the usual definition of convexity/concavity). With this definition in place we may make our main assumption regarding the profit function $\pi(f_m)$:

Assumption 3. The profit function is super-linear in imitation cost f_m .

This has several important implications. However, before discussing these we should ask whether this assumption is reasonable. Our answer is a clear yes. For example, in the Stackelberg case discussed above, and analyzed at length in Pollock (2006), profits are concave in imitation costs: $\pi(f_m) = \sqrt{f_m}$ (remember here that all values are normalized to lie in $0, 1$ so $\sqrt{f_m} \geq f_m$). Taking another of the examples mentioned previously, if f_m measures time before imitative entry then simple discounting again implies that π is concave and hence that π is super-linear. To take another example, suppose f_m represent the mark-up over (variable) costs the innovator is able to charge,¹¹ then for all ‘normal’ models of demand one again has that π is (quasi-) concave in f_m . Lastly, suppose the innovator has some fixed first mover advantage π_A over imitators so $\pi = \pi_m + \pi_A$. With a standard free-entry assumption imitator profits π_m will be driven down close to f_m . Then $\pi(f_m) \approx f_m + \pi_A$ and so π is super-linear in f_m .¹² To summarize, in any situation where profits are concave in imitation ‘costs’ π will be super-linear. In most normal cases (e.g. several of those just discussed) concavity follows from classic diminishing returns

¹¹This would be a natural outcome of a model where imitators locate in relation to an innovator in quality ‘space’ and f_m is some measure of their ‘proximity’ to the innovator – cf. ‘breadth’ in standard patent models.

¹²For $k \in (0, 1)$ we have $\pi(kf_m) \approx kf_m + \pi_A \geq k(f_m + \pi_A) = k\pi(f_m)$.

arguments and hence π is super-linear. We therefore feel this is a natural, and reasonable, assumption to impose. We now return to the implications of this assumption for the structure of the feasible set.

First, take any point in the feasible set (f_i, f_m) and $k \in (0, 1)$. Now, by super-linearity, $\pi(f_m) \geq f_i$ hence $\pi(kf_m) \geq k\pi(f_m) \geq kf_i$ and hence the point (kf_i, kf_m) is in the feasible set. In particular, for any point on the boundary of the feasible set (i.e. $(\pi(f_m), f_m)$ or, equivalently, $(f_i, g(f_i))$) the entire ray from that point to the origin must be in the feasible set. This immediately implies that $g(f_i)$, the function defining the boundary of the feasible set in terms of innovation cost, is sub-linear. Changing to f_i, ϕ coordinates this equates to the following proposition:

Proposition 4. *If an innovation (f_i, ϕ) is in the feasible set then all points on the horizontal line to its left lie within the feasible set. Equivalently, h , which defines the boundary of the feasible set, is non-decreasing in innovation cost.*

Proof. Take a point on the boundary of the feasible set: $(f_i^0, \phi^0), \phi^0 = h(f_i^0)$. Then $\pi(f_m^0) = \pi(\phi^0 f_i^0) \geq f_i^0$. Consider a horizontal line from this point to the ϕ axis, i.e. points of the form $(f_i^k, \phi), \phi = \phi^0, f_i^k = kf_i^0, k \in (0, 1)$. Then profits at a point on this line (f_i^k, ϕ^0) are:

$$\begin{aligned} \pi(f_m^k) &= \pi(\phi^0 kf_i^0) \\ &= \pi\left(\frac{f_m^0}{f_i^0} kf_i^0\right) \\ &= \pi(kf_m^0) \end{aligned}$$

By super-linearity of π , $\pi(kf_m^0) \geq k\pi(f_m^0)$ which we have already shown is $\geq kf_i^0 = f_i^k$. Hence (f_i^k, ϕ) is in the feasible set. Thus, the entire horizontal line is in the feasible set. This in turn implies that h is non-decreasing (if not we would have a point on the boundary of feasible set with not all of the associated horizontal line within the feasible set – contradiction). \square

The implications of this result, together with the previous remarks, may be summarized in the following proposition:

Proposition 5. *Take an innovation $X = (f_i, f_m) = (f_i, \phi)$ and another innovation $Y = (f'_i, f'_m)$. For some λ, μ we may write $f'_i = \lambda f_i, f'_m = \mu f_m$ and, without loss of generality, $\mu < 1$. If X is in the feasible set, a sufficient (though not necessary) condition for Y to be feasible is $\mu \geq \lambda$, that is, the drop in imitation cost is proportionally at least as great as the drop in innovation cost. In addition, if X is not in F and $\mu \geq \lambda$ then for μ sufficiently small (and hence λ sufficiently small) Y is feasible – that is if we reduce innovation cost and imitation costs at a proportional rate then for a sufficiently great reduction the innovation becomes feasible.*

Proof. X is feasible so $\pi(f_m) \geq f_i$. Now, $\pi(f'_m) = \pi(\mu f_m) = \mu \pi(f_m) \geq \mu f_i$. If $\mu \geq \lambda$ then $\mu f_i \geq \lambda f_i = f'_i$ and we are done. Note that a more intuitive, geometrical version, of this proof is provided using the IS diagram in the next section (see Figure 3). \square

3.2. Changing Regimes. We now consider the effect of a change of regime on whether an innovation is feasible – or more generally, for a given set of innovations what subset of them are feasible. In particular, we examine what happens if an open regime is adopted instead of a closed, proprietary, one. As discussed in Section 2.3, the change in regime is naturally modelled as a change in the innovation and imitation costs associated with an innovation. Thus, take a particular innovation or technology X , and define $X^C = (f_i^C, f_m^C)$ to be its location in innovation space under a closed regime and $X^O = (f_i^O, f_m^O)$ its location under an open regime. It will be useful to write X^O in terms of the coordinates of X^C by defining λ, μ, κ as:

$$\begin{aligned} f_i^O &= \lambda f_i^C \\ f_m^O &= \mu f_m^C \\ \phi^O &= \kappa \phi^C \implies \kappa = \frac{\mu}{\lambda} \end{aligned}$$

λ, μ , measure the *relative* drop in innovation and imitation costs (respectively) when moving from a closed to an open regime while κ is the ratio of these (relative) changes in imitation cost to innovation cost. A large κ correspond to a situation where innovation costs have fallen much more (relatively) than imitation costs, a small κ corresponds to the reverse, that is innovation costs have dropped (relatively) much less than imitation costs, and $\kappa = 1$ corresponds to the case where both costs have fallen relatively by the same

amount. Clearly we have $\lambda, \mu \geq 0$ as innovation and imitation costs cannot be negative. In addition we have:

Assumption 6. Under a move to an open regime imitation and innovation costs either fall or remain unchanged, that is $\lambda, \mu \leq 1$.

This has been stated as an ‘assumption’ however this statement appears to us so self-evident from the definition of openness as to hardly merit justification. Certainly, it seems clear that by making one’s work open one can only be reducing the costs of ‘imitation’ (after all reproduction and reuse are now expressly permitted), while it seems hard to think of any circumstance in which innovation costs increase (they may stay the same – more on this below).

Given this ‘assumption’, we always have $\lambda, \mu \in [0, 1]$ (and hence $\kappa \in [0, \infty]$) irrespective of the type of the innovation and its initial location under the closed regime. This provides one of the main reasons for using relative values rather than absolute ones (e.g. $\Delta f_i = f_i^C - f_i^O$), for, with them, we would have that the change in a variable such as innovation cost was limited by its initial value (e.g. $f_i \geq 0 \Rightarrow \Delta f_i \leq f_i^C$). Furthermore, relative values fit naturally with the previous results on the structure of the feasible set. We turn now, to the main question of this section. It should already be clear that the impact of a move to openness on feasibility will depend on how innovation and imitation costs change, that is on the values of λ, μ . We shall consider several distinct cases in turn.

3.2.1. *The ‘Traditional’ View.* The ‘traditional’ theoretical view discussed in the introduction equates to:

$$\text{Relative Change in Innovation Cost} = \lambda = 1$$

$$\text{Relative Change in Imitation Cost} = \mu \ll 1$$

That is, a move to openness, by depriving an innovator of the use of exclusive rights or secrecy,¹³ must necessarily reduce imitation costs – and reduce them significantly (hence $\mu \ll 1$ not simply $\mu < 1$). Furthermore, though usually more implicit, it is also assumed

¹³In fact secrecy is usually ignored and it is simply the lack of exclusive rights that is considered sufficient to justify this result.

that innovation costs are unaffected by a change a regime – that is, openness has no impact on the innovator’s cost of producing the initial innovation.

The implications of such assumptions are clear. Reducing imitation costs reduces an innovator’s profits while costs are unchanged. Thus, an innovation that is feasible under the closed regime may become infeasible under the open one (while any innovation that is initially infeasible remains so).¹⁴ Hence, moving to an open regime makes it less likely a given innovation is feasible and, thus, on average, reduces the level of innovation that occurs. Furthermore, with the drop in imitation costs being large ($\mu \ll 1$) it is very likely that an innovation feasible under the closed regime is infeasible under the open one.¹⁵

3.2.2. Changes in Innovation Cost. As discussed in the introduction there is accumulating evidence that an open regime can reduce the innovators cost either by simply making production more efficient or by allowing the innovator to share some of the costs with the wider user community. Let us therefore consider the more general case where a move to openness reduces innovation costs *in addition* to imitation costs:

$$\text{Relative Change in Innovation Cost} = \lambda < 1$$

$$\text{Relative Change in Imitation Cost} = \mu < 1$$

Here it will be useful to work primarily with $\lambda, \kappa = \frac{\mu}{\lambda}$ rather than λ, μ – that is in (f_i, ϕ) coordinates rather than (f_i, f_m) coordinates.

Proposition 7. *Irrespective both of the initial location of the innovation in IS and λ , if $\kappa \geq 1$ ($\mu \geq \lambda$) then an open regime is ‘better’ than a closed one. ‘Better’, here, meaning that the level of innovation under that regime is greater than or equal to that under the alternative.*

Proof. Proposition 5 shows that, irrespective of the initial costs and λ , if $\kappa \geq 1$ any feasible innovation remains feasible. Furthermore, it also shows that, at least for some values of

¹⁴In IS space we can visualize this change as move vertically downwards (note this trajectory is the same whether in (f_i, f_m) or (f_i, ϕ) coordinates). Since the feasible set lies above the line defined by a non-decreasing function g (or h) a move vertically downwards can only make it more likely a given point is not in the feasible set.

¹⁵Alternatively, suppose the drop in imitation costs is not large but (as is often also implicitly assumed) first-mover advantages are small then profits, $\pi(f_m)$, are close to imitation costs. Assuming that a given innovation is initially feasible we must have that its cost is close to (initial) imitation costs ($f_i \in [f_m, \pi(f_m)]$) and hence any reduction in profits is likely to make it infeasible.

λ (and/or the initial location of the innovation), a innovation infeasible under a closed regime becomes feasible under an open one. \square

This is the most general statement we can make without further information about the form of the feasible set (i.e. the structure of the underlying competition model) and/or details of the location of the innovation within that space. Nevertheless, it is important to realize that an open regime may be ‘better’ under much weaker conditions than $\kappa \geq 1$. For example, in the Stackelberg case discussed above $\pi(f_m) = \sqrt{f_m}$ and hence the boundary of the feasible set has the particularly simple form: $h(f_i) = f_i, (g(f_i) = f_i^2)$. Thus, in this case an open regime is better if $\kappa \geq \lambda$. So, for example, if innovation costs drop significantly under the move to an open regime, $\lambda \ll 1$, then the open regime will be better for all but very small κ (to be precise it will be better if imitation costs have fallen by less than $\mu = \lambda^2 \ll \lambda$ since $\lambda \ll 1$). To illustrate this point suppose, under the closed regime, that innovation costs were \$1 million imitation costs \$750k and that the innovation was feasible). Furthermore, suppose an open approach reduced innovation costs to \$200k, a fifth of their original value, then imitation costs could have gone down to any amount above \$30k (one 25th of their original value) and the innovation would still be feasible.

In addition to this ‘algebraic’ proof, it is useful to consider an alternative equivalent geometrical approach based on IS diagrams as this easier to understand and provides a more intuitive understanding of the results. In IS space using (f_i, ϕ) coordinates a change in costs given by $(\lambda, \kappa), \lambda \leq 1$, equates to a translation of a point in the following ways, each of which is illustrated on the IS diagrams in Figure 3:

- $(\lambda, \kappa) = (1, < 1)$: vertically downwards. This is the ‘traditional case’: innovation costs are unchanged and imitation costs ($\kappa = \mu$ with fixed innovation costs) have dropped.
- $(\lambda, \kappa) = (\lambda, 1)$: horizontally left. Innovation and imitation costs have dropped by equal proportions.
- $(\lambda, \kappa) = (\lambda, > 1)$: diagonally to the left and upwards. A reduction in innovation costs with a less than proportional drop in imitation costs. One (extreme) examples is the case of a reduction in innovation costs with no drop in imitation costs which corresponds to $(\lambda, \frac{1}{\lambda})$ (i.e. moving left along the hyperbola $1/x$).




FIGURE 3. Examples of different ‘trajectories’ in innovation space under a change from a closed to an open regime. Cases are numbered to correspond to those enumerated in text. Two different feasible sets are illustrated. The first illustrates perhaps the simplest profit function satisfying super-linearity viz $\pi(f_m) = kf_m$ (here $k = 2$). The second is that arising from the Stackelberg model.

- $(\lambda, \kappa) = (\lambda, < 1)$: diagonally to the left and downwards. A reduction in innovation costs with a more than proportional drop in imitation costs.

As just mentioned, Figure 3 illustrates these trajectories together with two examples of feasible sets both satisfying the (minimal) ‘horizontal-line’ condition. Not only is the basic result clear but we can easily deduce various other simple properties. For example, (for most feasible sets and most innovations) as $\kappa \rightarrow 0$ (imitation costs fall much more than innovation costs) a closed regime is ‘better’ than an open one (the ‘traditional’ result). Similarly, one can easily see how conditions on λ, κ will vary with the form of the feasible set (as given by its boundary function h) as well as with (initial) location of an innovation relative to that boundary.

To conclude, if we allow, as the evidence suggests, that openness results in a fall in innovation costs as well as imitation costs then an open regime may be superior to a closed one (for innovation). Whether this is so, depends on the relative rates at which costs of innovation have fallen relative to imitation costs. It was shown that, at a minimum, with equiproportional reductions in costs openness was always preferable (and therefore was also better for any situation in which innovation costs fell less rapidly relative to imitation costs). Thus, contrary to the ‘conventional wisdom’, openness need not always result in a decline in innovation – in fact it might sometimes increase it. Furthermore, as was illustrated by the use of an IS diagram, stronger results could be obtained if one imposed more structure on the feasible set and/or the initial location of the innovation under the closed regime.

On this last point, it is worth observing that producers of information goods may derive some benefit from their efforts, whether monetary or otherwise, that is not affected by imitation of the original work. For example, they may be able to supply an associated complementary item – e.g. support services – which are not open to easy imitation (perhaps they are rival or proprietary), or, alternatively, they may simply derive intrinsic satisfaction from their creative endeavours.¹⁶

The existence of such an alternative, but related, source of ‘income’, has three significant implications. First, and most obviously, it ensures a certain minimum level of ‘income’ even if openness has sharply reduced imitation costs and hence innovator’s profits from the

¹⁶See e.g. Rossi (2004) for a review of these sorts of possibilities.

original information good. Second, opening up the underlying information good is likely to increase usage thereby stimulating demand for the complementary one. Third, and least obvious, since knowledge of the original innovation is needed to supply the complementary one, ‘open-sourcing’, by increasing the gap between imitation and innovation costs, is likely to increase an innovator’s income on the complementary good by reducing the number of competitors.

Crucially, all three of these effects act a) to increase the size of the feasible set, and b) to increase it more at low levels of the imitation cost. In terms of the IS diagram this is an increase in the feasible set in the lower left-hand quadrant. Cursory examination of the IS diagram shows that it is precisely these sort of changes in the structure of the feasible set which make openness (more) attractive (e.g. compare top and bottom diagrams in Figure 3). They therefore make it more likely that an open regime is superior to a closed one.

4. CONCLUSION

In this paper we have presented a simple framework with which to under the interaction of innovation and imitation. Using it we discussed how a move from a closed to an open regime could be represented in ‘innovation space’ as well as how to represent (simultaneously) the set of feasible innovations – as determined by the underlying model of competition between ‘innovators’ and ‘imitators’.

It was shown that, when innovators enjoy first-mover advantages, innovation may still occur even when imitation is (substantially) cheaper than innovation. Going further we proposed a basic condition, which if satisfied by underlying model of competition (as encapsulated in the innovator’s profit function), allowed us to derive important features of the feasible set. With a regime determining the location of innovations in ‘innovation space’ one could use properties of the feasible set to make specific predictions about how changes in regime affected the level of innovation.

Starting from the accumulating anecdotal and empirical evidence that openness may confer significant efficiency advantages, we examined the situation where a move to open production resulted in a reduction in *both* innovation costs as well as in imitation costs – that is not simply a reduction in imitation costs alone as is usually assumed by the more traditional literature.

Precise sufficient conditions on the relative reductions in the two types of cost were derived under which a move to an open regime would actually increase innovation. These conditions had a particularly simple form: if the proportional reduction in innovation cost was at least as large as that in imitation cost then an open regime was better. We discussed the sufficiency aspect of this result and pointed out that in many situations it was likely that much stronger results would hold. Here however, though we provided some explicit examples to illustrate this point, we wished to remain at as general a level as possible and therefore did not go further in examining specific cases – this may be a fruitful area for future work.

Nevertheless, we did observe that the ability of an innovator to sell some form of complementary good would be one major example where stronger results could be obtained. A cursory examination of current ‘real-world’ practice reveals that this possibility is more the norm than the exception. For example, the author of a piece of software is often able to supply support services or additional ‘plugins’. Similarly a musician gives live performances in addition to writing the music itself, and engineers may charge for consultancy as well as access to any inventions they make.¹⁷

Thus, it is likely that in a variety of areas open approaches are a sustainable form of operation and, at least in some cases, may be superior to proprietary ones. However, as should be clear, we are not suggesting that this will always be the case. It is therefore important for society to have some method for deciding for which industries, or technologies, it is so, and for which it is not. One of the primary contributions of this paper has been precisely to provide a simple, intuitive, yet rich, framework in which, by reference to an overarching ‘innovation space’, one can go about doing this.

REFERENCES

- Arundel, A. (2001). Patents in the Knowledge-Based Economy. *Beleidstudies Technology Economie*, 37:67–88.

¹⁷This last case is inspired by the 18th century engineer John Rennie, one of the most famous engineers of the industrial revolution. In 1789 he worked on the Albion Mills for Watt and Boulton. To Watt’s horror, upon completion, Rennie, rather than patenting his new design, was eager to demonstrate it to others. “[F]ar from ruining him [Rennie] as Watt predicted, [this] established his reputation and led to a flood of commissions” (Macleod, 1988, p. 104).

- Bessen, J. (2006). Open Source Software: Free Provision of Complex Public Goods. In Bitzer, J. and Schraeder, P. J. H., editors, *The Economics of Open Source Software Development*. Elsevier B. V.
- Cohen, W., Nelson, R., and Walsh, P. (2000). Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not). NBER Working Paper No. W7552.
- Hopenhayn, H. A. and Mitchell, M. F. (2001). Innovation Variety and Patent Breadth. *The RAND Journal of Economics*, 32(1):152–166.
- Klemperer, P. (1990). How Broad Should the Scope of Patent Protection Be? *RAND Journal of Economics*, 21(1):113–130.
- Koch, S. (2008). Effort Modeling and Programmer Participation in Open Source Software Projects. *Information Economics and Policy*.
- Lakhani, K. R. and von Hippel, E. (2003). How open source software works: "free" user-to-user assistance. *Research Policy*, 32(6):923–943.
- Levin, R., Klevorick, A., Nelson, R., Winter, S., Gilbert, R., and Griliches, Z. (1987). Appropriating the Returns from Industrial Research and Development. *Brookings Papers on Economic Activity*, 3:783–831.
- Macleod, C. (1988). *Inventing the Industrial Revolution: The English Patent System, 1660-1800*. Cambridge.
- Mansfield, E. (1985). How Rapidly Does New Industrial Technology Leak Out? *Journal of Industrial Economics*, 34(2):217–223.
- Nordhaus, W. (1969). *Invention, Growth and Welfare: A Theoretical Treatment of Technological Change*. M.I.T. Press.
- Pollock, R. (2006). The Value of the Public Domain. Published by the Institute for Public Policy Research as part of a series on IP and the Public Sphere.
- Rossi, M. A. (2004). Decoding the "Free/Open Source(F/OSS) Software Puzzle" a survey of theoretical and empirical contributions. Department of Economics University of Siena 424, Department of Economics, University of Siena.
- Scherer, F. (1972). Nordhaus' Theory of Optimal Patent Life: A Geometric Reinterpretation. *American Economic Review*, 62(3):422–427.

Scotchmer, S. and Green, J. (1990). Novelty and Disclosure in Patent Law. *The RAND Journal of Economics*, 21(1):131–146.

Teece, D. J. (1986). Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research Policy*, 15(6):285–305.