Hobbyists and the appearance of new industries

by Peter B. Meyer, of the Office of Productivity and Technology
U.S. Bureau of Labor Statistics
June 6, 2004
Preliminary and incomplete

Abstract. Some inventions with great economic impact were produced by networks of developers working together. The airplane, the personal computers, and various open source software projects are examples. The networks existed prior to the inventions and seem to have been essential to the success of the eventual inventors. This paper attempts to model the rise of the network. It defines a search space for a low cost technology that would have value to customers. When the technology is advanced, the agents have incentives to compete. When the technology is primitive, they have little incentive to compete but may not wish to share. The paper considers various utility functions to see which ones can describe creators of the network. A hobbyist utility function for example is defined to describe people who have an independent interest in experimenting with a particular technology and who may want to become entrepreneurs. Simulations suggest which types and parameter values motivate players to choose to share their technological discoveries. The model provides a way to think about which societies and economic environments are more or less supportive of collective invention.

1.0 Introduction

No individual just invented the airplane, the personal computer, or Linux. These inventions arose through a process in which many people were trying to make something new. Many of them shared technological information through clubs, journals, or professional societies. It would have taken longer for these important technologies to appear if the information had been kept secret. This process by which technology advances are freely shared has been called *collective invention*.[1] When software source code is shared, revised, and shared again (the open-source software approach), the same kind of process is operating, in an especially efficient environment.

Individuals who choose to reveal their technology to others may feel tension about this choice since if instead they kept their technology secret, it might have become valuable private information if and when there were to become great demand and an

---

[1] The term originally comes from Allen (1983). This paper uses the broader definition of Meyer (2003): Collective invention is a *process in which improvements or experimental findings about a production process or tool are regularly shared.* It does not mean all investigators are sharing, but some regularly do.

industry of suppliers appeared. Historically the collective invention process has sometimes broken down after strongly profitable opportunities appear from the technology.

In following the story line in this paper we leave out whether or not the inventions were formally intellectual property protected by copyright or patent – the core point relevant here is whether the invention was in the end made available. A patentholder could thus make either choice in the collective invention context, either licensing the invention easily and widely, or keeping it unavailable.

The purpose of this paper is to show that an abstract model of hobbyists, behaving optimally, can exhibit collective invention behavior. Such a model gives us a way to make predictions about collective invention in the context of different environments and counterfactuals. For example, the model gives us a way to think about what environments and actions by governments can prevent collective invention from occurring. During the Industrial Revolution, Britain seems to have been a good place for collective invention. Put differently, "the key to British technological success was that it had a *comparative* advantage in *micro*inventions." (Mokyr, 1993, p. 33) In the past two centuries, new technologies spawn new industries most quickly in the United States. Collective invention may explain why.

Collective invention may disappear over time in the model as an internally competitive industry arises. The model includes (a) a utility function for hobbyists, according to which they are interested in improving the technology whether or not they later become entrepreneurs or are personally rewarded; (b) a space of possible technologies defined by two variables -- cost and production capacity – along which the agents have slowly improving technologies which they may choose to share; and (c) shared technology advances have the effect of causing standardization in a path-dependent way. It is hard or impossible to solve the model analytically, but the expected-utility-maximizing choices are computable at each technology location, which makes it possible to simulate histories in which hobbyists begin work, and eventually a profitable industry forms.

In the next section, several episodes of collective invention are summarized to establish that the phenomenon is important although not common. After that the model and its computable solution is described, and the resulting simulated technological histories discussed. Then in the last section various implications are highlighted.

2.0 Collective invention histories

The stylized accounts that follow illustrate the phenomenon which is to be modeled. The compressed histories focus on particularly important inventions both because this is especially interesting, and because it is in the case of really important inventions that the technological uncertainty is greatest and therefore the collective invention phenomenon is most important. The cases to follow are: steam engines, mass produced steel, automobiles, airplanes, personal computers, and open-source software projects generically.

2.1 Steam engines

Before there were steam engines, there were a number of people who foresaw that there might be some way to use fuel to get mechanical work done by boiling water and applying the resulting steam pressure. In Robert Thurston's account of this, operators of mines focused on this problem because they had easy access to wood for fuel, and they wanted to pump the water out of mines so they could dig deeper. By 1675, the problem and the opportunity were recognized by Edward Somerset, who tried to make such a device but never succeeded. In 1698 Thomas Savery patented a kind of steam pump which could pump water out of a mine if enough wood fuel were burned, but it was so inefficient that it was not used much. Later accounts say that it worked in an engineering sense but was not economical to use. In or around 1711, Thomas Newcomen invented what is now described as the first real steam engine. It was rapidly improved by a number of people and was more or less standardized by 1718. It is not clear that any of the people meant to go into the business of producing steam engines – rather, they wanted to *use* a steam engine for their application. Hundreds of Newcomen engines were used in British mines by the 1750s.

James Watt formulated a design for a compound steam engine in 1765 and correctly foresaw that it was a dramatically more efficient design than the Newcomen engine. He could not find tools precise enough to make the necessary pistons for several years until he worked with machine maker John Wilkinson who designed a special machine to bore out smooth cylinders. Watt also worked closely with a business partner who decided they should hold the patent tightly. Some steam engine engineers copied the design before the patent expired. After 1800, when the patent did expire, steam engine owners shared technological information through the publication called *Lean's Engine Reporter* (Nuvolari, 2001). Much of the inventive work on steam engines both before and after Newcomen's first steam engine, and in the later period after the expiration of Watt's patent, fits in the category of collective invention.

Iron work

Iron producers in Britain from the 1850s to the 1870s, and steel producers in the United States starting in the 1860s, shared technological information regularly with one another. (Allen 1983, Meyer 2003)

Airplanes

For a hundred years there were tinkerers trying to figure out how to make powered vehicles passengers safely through the air. Many hobbyists, engineers, and other specialists devoted energy to aeronautics although there were many people who thought it was unrealistic. The story is told elegantly and in detail in Couch (1989). Some of the specific steps forward are distilled in table A2.

Personal computers

Microcomputers were first designed, built, and modified by hundreds of hobbyists and tiny startup companies who served them. The most famous hobbyists met monthly at the Homebrew Computer Club near Stanford University starting in 1975. Members of the club generally shared technology freely with one another, sometimes for ideological reasons. Apple Computer and many other such companies spun off from work done by these hobbyists. (Meyer, 2003; Levy, 1984; Freiberger and Swaine, 1984) Some of their inventions are listed in Table A3.

Open source software

Free or open source software projects are those on which the source code is more or less freely shared with other developers for them to use or improve. The argument that this a collective invention process has been made by Nuvolari (2001).

Stripped of technological and institutional detail, we can see a general shape to these histories. It begins when some kind of technological opportunity is perceived to appear and interested parties try to advance the technology for a variety of reasons of their own. Some choose to share their advances with others. Eventually there can be an advance that makes a profitable industry profitable, in which case the collective invention process is likely to break down because the incentives to participate have weakened so much.

3.0 Why they shared and standardized

People experiment with and share technology for a variety of reasons. Some are purely private benefits, and others come about through a change in the environment, most of which is a public benefit. This partial list is rich enough to show why making a detailed utility function for hobbyist experimenters is a complex problem.

- For the “use value” – because they have a use for the resulting creation

The first steam engine makers did this – it is not clear that they expected to be able to manufacture and sell steam engines, but it is clear that they meant to USE steam engines, to pump water out of mines.

- fun, excitement
- prestige, recognition, fame (for either individual or employer)
- employer’s own technology might improve through paybacks, alliances, or side effects
- impractical or costly to keep secrets, e.g. because of high job turnover
- better public technology raises value of assets owned by the innovator

The technological change also produces indirect effects on the person or institution which shared the technology. If the new technology is useful to others, a firm’s release of information about it could increase the value of some asset it had. Robert Allen (1983) highlighted this effect among British iron producers of the 1850s to 1870s. When better ironmaking methods were introduced, ore deposits owned by a British iron firm gained value, whether or not that firm itself used the new methods.

- To establish desirable engineering standards (for market power, perhaps, or to make a particular feature universal.)

Netscape and Microsoft did this when they improved their browsers then gave them away in the late 1990s.

- For the employer to evaluate employees (as in "open science" or academia)

An individual can get some private benefits unrelated to changing the technology through paths like these. First, hobbyists and people within firms have ambitions for prestige, fame, and future employment opportunities that they can help meet by releasing technical information. Second, firms garner publicity by making their successes known, and such announcements may encourage their staff. Third, an organization may not intend for information to be released, but does not find it worthwhile to spend the costs necessary to keep it secret. This is especially the case when there is substantial movement of employees between firms. Fourth, publications in an open environment give employers a way to judge the contribution or skills of a researcher which may be hard to judge directly. Publications approved by editors serve as a certification signal as suggested by Paul David (1998) in a scientific, not technological, context. These benefits operate through channels unrelated to real effects of the technology change. Some economic researchers (such as Lerner and Tirole, 2002) explain sharing behavior solely in these terms, although it does not accord well with the descriptions of the technologists themselves.

There are also public-benefit reasons:

- Fifth, a firm may expect that cooperating with other firms to improve their production process would induce improvements to its own technology, either as a side effect or as a payback by the recipient firm. In computer and software contexts especially there are advantages from establishing engineering standards by giving away designs or software. This was one reason Web browsers were given away.
- Sixth, although each firm competed against other local firms, collectively they compete against other regions. They have an incentive to work together to make local production as efficient as possible and the remote regions irrelevant. They might do this by allowing consultants to consult with other firms as well as their own, by encouraging local suppliers to gain economies of scale, by building a common transportation infrastructure, or by agreeing to engineering standards.

Those incentives apply to firms. Hobbyists may have different reasons. For them,

- Playing with the new technology can be fun and absorbing, and many computer hobbyists thought it was virtuous[2] and exciting to share their findings.

Some technologists may have a particular need they are trying to satisfy. Programmers may be willing to write a fix or feature in software because of a particular itch the programmer wishes to scratch. (Raymond …)

Why do they SHARE?

- virtuous

Stallman. Raymond? Lakhani?

- employment opportunities

As in Lerner and Tirole, and the FOSS-economics discussion generally.

- to speed development, as discussed in Foster (1986). His discussion of the IBM PC says IBM shared tech in order to move faster. This was a way of operationalizing their decision to spend more-than-usual-resources to achieve faster-than-normal product development.

The force at work. There is a direct meaning to making history here in the model which has a real world analog. (or: the model reflects a real-world economic force that is not of the career-concerns type). If I give the other player my technology, then his innovations are more likely to be directly useful to me In the open-source software case, his bug fixes are then useful to me, and he may share them with me directly. There can be indirect benefits too. A buggy sendmail in China may have costs to its author in Pennsylvania and he would prefer to upgrade the user there fore free than for the older, incompatible or buggy version to be in use anywhere. In the case of technologies that require physical presence and tinkering, (e.g. steam engines, automobiles, and airplanes), it is also the case that one might want to have with the others common tools, materials, skills, and so forth.

One has been slighted in the formal economic literature on open-source software, and the model to be shown is built around that reason so as to verify that it is potentially sufficient to explain the development of the technology and the rise of an industry in the model.

Another economic force at work is that there can be a payoff to changing the world around the agent, apart from changing his own position within it. The force at work here is not principally career concerns, or the desire for recognition and prestige. The impulse is to improve the situation for the other person, literally so that it is better for him, with ancillary benefits to oneself. For example, one may not want to hear a colleague complain about the software she is using, or about a squeaky door, and therefore one has the impulse to fix the door and upgrade her software.

---

[2] Many innovators express this thought in the descriptive literature. They say they are doing what is *right* (because software *should* be shared, "like recipes" in Stallman's language) or are trying to achieve something *good* (because society can be made better through this kind of contribution). This view does not usually conflict with paid employment or secrecy in other parts of their lives.

There is a direct effect of making history by improving the situation permanently. There is also a benefit of standardization, which is that the other player's innovations would down the road be useful to oneself.

One might want one's colleague not tot have a squeaky door. Thus one might directly wish to improve his technology, not because of career concerns, but to make one's own world a better place. For the same reason, one might prefer that other people did not complain about the food, water, shelter, or other environmental conditions they deal with, and this is one incentive to improve them.

In OSS we see the same phenomenon within each project. Let us focus on the biggest, best-known project, which makes the operating system Linux. If a developer has made a private improvement to the operating system, and later upgrades to a later version of the operating system from some outside source, he will either lose the benefits of his earlier improvement, or will have to fold his own improvement together with the new version, at an extra cost of time, energy, and the risk of making a mistake. This is one fo the costs of "forking" in project histories, that is, of divergent standards in product engineering. The problem of folding two versions back together to make a product with the best features of both is not solved only once, then, but must be done against and again whenever one upgrades. The usual and preferred solution, which engineers learn to treat as natural, is to unify the improvements into a single version if possible so one never faces the forking problem more than once with each set of innovations. The "owners" of an open-source project may try to achieve this. It is a technological and also interpersonal, political, challenge. If they can avoid forking, this is one of the benefits of standardization. (cite Info Rules) (and/or Everett Rogers)

The same force applies in the previous episodes and for clarify let us highlight some examples. If one is using an early steam engine – not selling, but using – it may be helpful if one's neighbor funs exactly the same kind, so that the two of you might share a stock of spare parts, or share repair skills, or loan tools to one another. Analogously if one is using a personal computer it is helpful if one's neighbor is using a compatible one so that one can share skills and software with the neighbor.

So there may be both a direct effect of upgrading the other player's technology, and there are network benefits from standardization. Only the second benefit is in the model. The other is too straightforward to require modeling, and introducing direct benefits would also in the literature raise the issue of direct costs.

[somewhere in here, cite Hackers, Cringely, and Raymond. Also Shapiro and Varian]

The model below treats this economic impulse – standardization on the latest technology – as the central force which makes the players want to share technology. While the other forces exist, we abstract from them here to establish the importance of this one. We model the hobbyist as having this utility – like that of a firm, but with a personal desire to do this particular activity. Each hobbyist therefore makes a profit or loss, computed by total revenues minus total costs, each period. This profit is ($p_t q_{it}$ –

$c_{it}q_{it}$). Let the per-period utility for a hobbyist be the profit plus a subsidy the hobbyist is willing to give each period. This is the resulting *hobbyist utility function*:

$$u(i,t) = s_i + (p_t q_{it} - c_{it} q_{it}) \qquad (1)$$

The subsidy is a large positive constant for each hobbyist. Hobbyists are willing to lose some fixed finite amount of money each period to apply this particular technology. In the real world analogs, the subsidy might come from the benefits this technology provides to other work that the hobbyist does – for example, the early steam engine makers did not plan to sell steam engines, ever. Rather, they wanted a steam engine in order to use it. Another reason we might see a subsidy is that the hobbyist enjoys working with the relevant tools and facing the challenge of making the technology, such as a primitive computer, work. A third kind of subsidy would come from professional investigators like scientists or academic research engineers, whose research brings them prestige, keeps them up to date, and, in centuries past, patrons. The first kind of subsidy comes from the use value of the technology, the second we could call enjoyment, and the third prestige value. These distinctions are often clear historically but do not affect the model. The key summary abstraction here is that these particular agents are willing to subsidize apparently unprofitable use of the technology.[3] This is not meant to describe people in general, but only the few who are relevant to the early history of a radically new technology.

[[on hobbyist utility: Imagine that for this particular field of activity, all individuals are arrayed along a line that describes how much they have to be paid to participate in the activity. Most people would be willing to work in the computer business as long as they got a reasonable salary. But those people are irrelevant to the invention of the personal computer. We won't model them. Think about people at the other extreme. The people who invented the personal computer are the kind of people who would be willing to pay $2000 a year to work on computers. Or in the airplanes case, the kind of people who would be willing to build their own hang glider out of wood, then strap themselves into it and jump bodily off a mountain.

Call each person's type s, for subsidy. Most people are willing to give only a negative subsidy to the computer

---

[3] Perhaps this utility function would describe other voluntary activities, and perhaps it has been done although I am not aware of it. Examples include the behavior of authors, scientists, volunteers, fanatics, evangelists, revolutionaries, and squatters and their associations (de Soto 2000, pp 137-147). Artists, designers, or musicians starting a new movement or style face a similar problem and could be modeled with a similar utility function. The cognitive experience between these categories of volunteers may be similar across these kinds of people. They want to change the world. They may also want to get credit for changing the world, but that sort of payoff is meant to be included in the profit expression. Here, we are looking at something different. They want the world to be different and are willing to subsidize the effort to make it happen.

industry. Our model will focus only on the two people who are willing to give the most extremely positive subsidy. But it will become clear how the others get integrated.

These special individuals would like to be in the business of computers or airplanes or whatever. But if there is no business, they'll experiment a little anyway because they have some private reason. (We saw a whole list of possible private reasons in the last section. The details are irrelevant -- people have tried to puzzle out the relationship between creative acts and where they come from, but I think that's beside the point. In a population of millions, there's going to be one who is willing to jump off the mountain. Done.]]

[[possible note on why we might think that the creators have a special type. After the telephone, Alexander Graham Bell was drawn into aeronautics (Crouch). Holley was in rails and steel. Radio culture lasted into west coast transistors. Did chip guys get in on the web?]]

3.1 The new technology stage

When we say that the hobbyist has a particular interest and particular skills, how do we define the activity? It is not yet an industry. We envision it as a search process for a technology that the player believes will exist. The service to be provided by the technology is, for example, controlled flight through the air, or interactions with a computer that one could reprogram or rebuild or add features to, or an operating system that one could understand entirely from the iniside and out, and which no one else would have property rights over. So let us assume the goal is well envisioned by the hobbyists, and that means exist to achieve this goal, poorly.

Let c be the quality-adjusted cost of providing a certain service, per unit produced. Let q be a quality-adjusted quantity of the services that the hobbyist can provide in the period. These two attributes describe the technology a hobbyist has available at any particular instant. We will be following the dynamics as our hobbyists move through this technology space, depicted in the figure below. The horizontal axis measures cost, transformed to be (1/cost) so that production is infeasible along the vertical axis, and production along the horizontal axis is not meaningful. Actual production can occur when the hobbyist has a technology above the horizontal axis and to the right of the vertical axis. Technological progress represents movements of technologies rightward and upward. These will occur over time in the model.

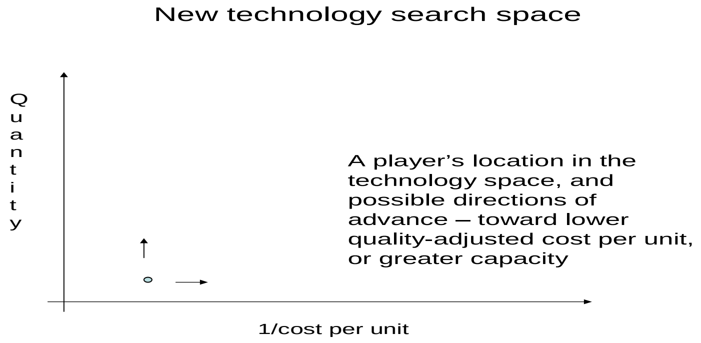

In this model the flexibility implied by a production function or cost function is imagined not to exist. The hobbyist has the option of producing quantity q, at cost per unit c, or not doing so.

There is however a demand function, taken to be exogenous. Let p be the quality-adjusted real price per unit of this service which a market that includes the hobbyist is willing to pay, given a quantity that is made available..

3.2 Simulation environment

To get the model to close down it has to end. Let there be an outer edge of levels of c and q beyond which technological progress is not possible. Eventually the players progress to the lowest cost technology possible.

We fill in a matrix in which every combination of (c1,q1,c2,q2) from (.1,.1,.1,.1) to (4.0, 4.0, 4.0, 4.0) is a grid point, and the present discounted value of arriving at the point are computed, and the optimal production, adoption, and sharing decisions are computed. We begin to fill the matrix at the outer edge where it is computable. Then, one notch in from the edge the values and optimal choices are computed because subsequent payoffs are known. Continuing, we arrive by backwards induction at the beginning. Having changed some of the payoffs passing through, we re-pass it a few times till it converges.

4.0 Model and method

4.1 Hobbyist utility and profit

The model has two agents (i=1,2) with what we'll call a *hobbyist utility function*. Each period (t=1,2,3, …), the hobbyists produce a service which they can use or sell. Each hobbyist has a technology to produce the service. The technology is characterized by (c_{it}, q_{it}), where c_{it} is a quality-adjusted cost per unit of the service produced, and q_{it} is a quality-adjusted quantity of services that would be provided. By quality-adjusted we mean the cost and quantity are expressed in service units that are comparable over time.

There is a market for the services. Total quantity produced in a period, t, is

$$Q_t = q_{1t} + q_{2t}$$

The market price p_t for the services is determined by this inverse demand function, known to the agents, in which k is a fixed positive parameter.

$$p_t = k/Q_t$$

Each hobbyist therefore makes a profit or loss, computed by total revenues minus total costs, each period. This profit is ($p_t q_{it} – c_{it} q_{it}$). Let the per-period utility for a hobbyist be the profit plus a subsidy the hobbyist is willing to give each period, as discussed above:

$$u(i,t) = s_i + (p_t q_{it} – c_{it} q_{it})$$

The subsidy is a large positive constant for each hobbyist. Hobbyists are willing to lose some fixed finite amount of money each period to apply this particular technology.

Each hobbyist makes choices to maximize their present discounted utility over all remaining periods, meaning up to some end point T which the hobbyist is presumed to know. Future utility is discounted by ß each period ahead, where ß is a constant near but below one.

$$U(i,t) = u(i,t) + ßu(i,t+1) + ß^2u(i,t+2) + ß^3u(i,t+2) + ... + ß^{T-t}u(i,T)$$

At any time an agent could cease production and receive a utility of zero in that period. Since the hobbyist could do this forever, the hobbyist need never accept a value of U(i,t) which is negative; the outside option worth U(i,t)=0 is always available. By these rules, an agent might accept a per-period negative utility u(i,t)<0 if future payoffs were high, although this case is not important for the model under consideration.

4.2 The stage game

Each period a series of decisions occurs. First, each player decides whether to produce the good. Quantities are determined by the technologies held by each player, and prices are endogenous, determined by an unchanging demand function.

As a side effect of production, one player may have a technological insight, which we model as an offer of a technological improvement from Nature. The improvement is either a rise in the quantity capacity technology, or a reduction in the per-unit cost.

The player who just received such information may use it or not, and if using it, may share it with the other player. The other player may accept it and if so jumps to the offered level of costs or quantity. The model does not include institutional detail but we could imagine that the sharing process goes through a technical journal or a club.

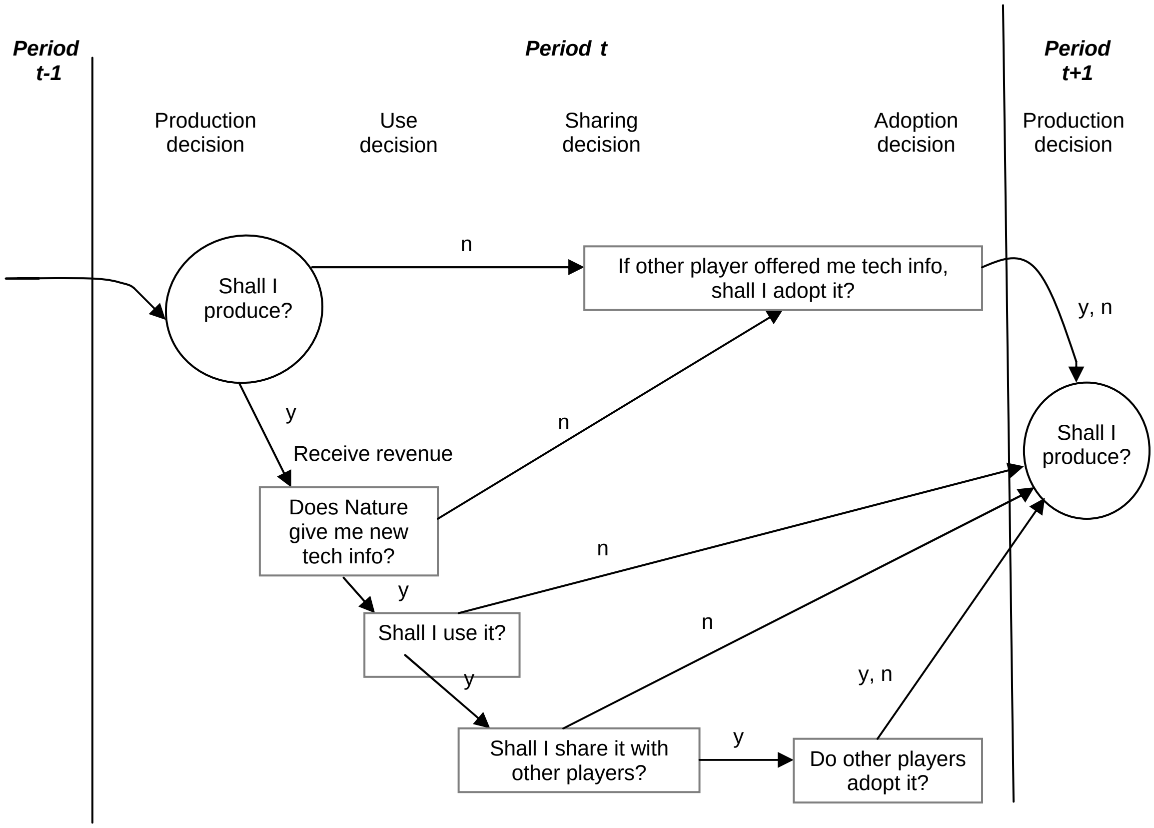

Figure 1. The stage game

4.3 The sharing conjecture

We wish to show that there exist parameters for which players choose to adopt and to share certain kinds of innovative discoveries early in the game.

First let us constrain the case being considered here to those cases where both players are willing to participate from the beginning. Assume they start (at time t=1) at a location with primitive technology. For an illustration, let both players start with technology of (c=5, q=.2) at time 1.

We will assume that the s_i parameters are such that the immediate utility in the first period for each of the players is high enough that they are willing to produce the good in the first round, even if it had no future payoffs:

$u(1,1)>0$ and $u(2,1)>0$

It follows that:

$s_1 > c_{11}q_{11} - q_{11}k/(q_{11}+q_{21})$ and $s_2 > c_{21}q_{21} - q_{21}k/(q_{11}+q_{21})$

For any $k>0$, such values of s_1 and s_2 exist.

Table 1. Parameters restrictions

Parameter restrictions in the model	**Parameters in the simulation**
$k>1$	
$ß>.5$, $ß<1$	
$s_1 > c_{11}q_{11} - q_{11}k/(q_{11}+q_{21})$	
$s_2 > c_{21}q_{21} - q_{21}k/(q_{11}+q_{21})$	

The agent’s choice about sharing is the result of counting up various costs and benefits. One of the costs of giving away the technology is that if the other player takes it, it is easier for him to advance to the industrial, competitive stage. One of the benefits is that when the other player has the same technology, any technological advance the other player is also an advance to the first player. This is a benefit from standardization, and occurs repeatedly, in the same sense that an organization may repeatedly standardize on new versions of software whether it makes them or buys them. Within the environment specified agents make a tradeoff between these costs and benefits.

The sharing conjecture is that for *some parameters, at some stage, the optimal decision for a player is to share his technological innovations freely*. By discussing when this is the optimal behavior for a player in the game we hope to develop some insight into when collective invention occurs in real life.

Indeed in simulations one finds that if the grid includes no spaces where the technology is profitable but there is some scope for the technology to cost less and some time for the players to benefit from one anothers’ advances, the players do share, and advance together.

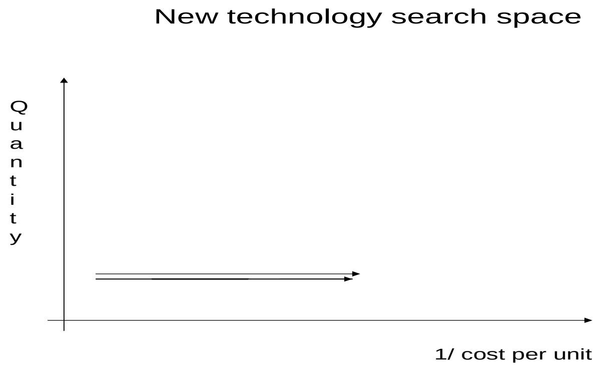

Figure 2. Players with technology that will never be profitable advance together

The model of Saint-Paul (2003) has the same incentive, shown occurring among profitable firms within a competitive industry. There, an equilibrium exists in which firms share technological information. A Nash equilibrium exists in which they do not share information, too, just as it would here if players considered one another's strategies in advance. Saint-Paul's model does not have the self-destructive property here; either equilibrium is permanent. But it is apparently tractable analytically, with value functions for the present-discounted values of alternative choices, and true proofs are discussed.

Cite also harhoff, henkel, and von hippel (2002).

We have seen then that if the technology is always unprofitable, there exist parameters for which the players share endogenously.

4.4 The backwards induction problem

Whenever grid is wide enough to cover profitable locations, a backwards induction problem appears which prevents collective invention from appearing. Just before a technology becomes profitable, players do not wish to share. Since in the previous round each foresees that the other will not share in the next round, it is not optimal to share previously either. This argument can be applied right back to the beginning. Neither is willing to be the last to share, therefore a sharing strategy never seems optimal, even far in advance. This occurs because of the artifact that there is no uncertainty in the model – the players know the probabilities and the possible technologies, perfectly. In real life, there are many things such players do not know; for example, they probably do not know with certainty that the technology will ever be profitable. Small perturbations in the model in realistic directions will make sharing optimal. We might assume either that (a) players get a small extra payoff ε from sharing, or (b) that new technology sometimes leaks out, and is shared even when the innovator did not mean to share it, [or maybe (c) they may receive a setback which a public tech helps them recover from].

4.7 Findings so far

For certain parameter values, early in the game, the players choose to share cost information but not quantity information since they both benefit by lowering their costs, and neither wishes to produce more than they have to in order to stay in the game, since each marginal unit produced loses money for them. This then is a theorem.

Later in the game there is a break point where they start to compete, by accepting quantity information from Nature. If there is any prospect of driving the other one out of the game or to force the other one not to rather to behave as an industrial competitor but instead to continue in a hobbyist mode, they also cut off sharing cost information. Thus, the hobbyist club has become an industry.

The presence of more hobbyist-firms makes the force stronger – that is, the incentive to share is stronger, as in Saint-Paul (2003).

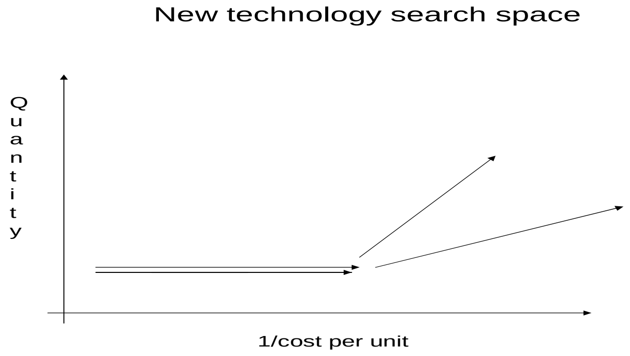

5.0 Implications for interpretation and policy

A social environment permitting a free press and free assembly enables this process to occur. If these things were not permitted the process would slow or stop. This gives us a partial explanation of why societies which forbid these things do not invent the greatest technological advances. is good. Cheap publishing technology is good – as in, first, the moveable type printing press, and now the Web.

There is a historical aspect of the model which could be misinterpreted. One need not believe that for every technological puzzle there exist experimenters willing to subsidize the process. But when there are, it can be very important, because they can move the process forward much more quickly than it otherwise would go. The alternative to progress through collective invention may be no progress at all. And when that process makes a difference, it can have a permanent, non-ergodic effect on technological progress and economic growth.

Why have such a model? One reason is that in economics, evidence must be defended by theory. Less archly, we can see that some economic discussions will treat it as irrelevant until there is a theory of it. For example, Lerner and Tirole (2002) . . . .

Some policy and environmental variation can be imagined here through the model. Suppose freedom of the press or freedom of assembly were not allowed. Then the players would have more difficulty finding others with the same interests and sharing at will. Working alone they would have more difficulty making progress and so invention would slow. If publication is expensive there the model would predict less sharing. In a world with the Web then there should be more sharing. If the society is rich, relative to the costs of innovation, there would be more collective innovation because it would be easier for the players to subsidize experimentation. If there is a common language (e.g. English) collective invention is easier. If the subject draws interest and fascination and they (macrosociologically) *believe* in it, there would be more collective invention.

6.0 Conclusion

This paper has displayed a model economic environment and a utility function for a hobbyist whose optimal behavior in the model produces collective invention and the transition to a competitive industry. The transition in the model has analogues to transitions historically observed in the histories of aircraft, personal computers, and other radically new inventions.

7.0 Appendices

Table A1 Relevant innovators in the proto-automobile industry

Automobiles

date	who	advance	source
1789	Nicolas Joseph Cugnot	first self-propelled vehicle -- steam-propelled, with three wheels. 3 miles/hr	Columbia, p.194
1801	Richard Trevithick	another three-wheeled steam-driven car	Columbia, p.194
1885	Karl Benz	three-wheeled car with internal-combustion engine	Columbia, p.194
1885	Gottlieb Daimler	three-wheeled car with internal-combustion engine	Columbia, p.194
1890s	many manufacturers in US	modern-type internal-combustion cars	Columbia, p.194
1892	Henry Ford	built his first car	Columbia, p.977
1894	Stephen Balzer	air-cooled three-cylinder internal combustion engine. New York City.	Crouch, p. 260
1895	George Selden	patent. Patent pool started 1903.	Columbia, p.194
1897	Stanley brothers	produced Stanley Steamers, steam-driven cars, in Massachusetts	Columbia, p.194
1903	Henry Ford	co-founded Ford Motor	Columbia, p.977
1908	Ford Motor	Model T Ford	Columbia, p.977
1914	Ford Motor	$5 for 8-hr day minimum wage; also profit-sharing plan	Columbia, p.978

Table A2 Relevant innovators before the airplane

Airplanes

date	who	advance	source
1804	George Cayley	tested flying model gliders, which had "cruciform tail" combining rudder and elevator; formulated fundamental problem of aeronautics (how to apply resistance of air)	Crouch, p.28
1823	A.A. Mason	failed Aerial Steamboat	
1858-9	F.H.Wenham	glider with "cambered" wings	
1866	British engineers	founded Aeronautical Society of Great Britain	Crouch, p. 30
1868	Britain	beginning of publication of Annual Reports of Aeronautical Society	Crouch, p. 31
1869	Paris	beginning of publication of *L'Aeronaut*	Crouch, p. 31
1871	F.H.Wenham and John Browning	developed wind tunnel	Crouch, p. 31
1875	Octave Chanute	discovers, on trip to Europe, that European engineers treat airplane as real possibility	Crouch, p. 26
1888	France	beginning of publication of the *Revue de l'Aeronautique*	Crouch, p. 31
1903	Wilbur and Orville Wright	self-powered flight, with takeoff and landing at same level	
1904	Wrights	testing grounds on Huffman Prarie	
1911	Italy	First use of airplane in war by Italy against the Turks in Tripoli	

Table A3 Relevant innovators before the standardized personal computer

8.0 References / bibliography

Allen, Robert C. (1983) "Collective invention." *Journal of Economic Behavior and Organization* 4: 1-24.

Cringely, Robert X. (1992), *Accidental Empires: how the boys of Silicon Valley make their millions, battle foreign competition, and still can't get a date*. (edition used here: HarperCollins, 1996.)

Crouch, Tom D. 1989. *A Dream of Wings: Americans and the airplane, 1875-1905*, second edition. Norton.

David, Paul. 2000. "Patronage, Reputation, and Common Agency Contracting in the Scientific Revolution: From Keeping 'Nature's Secrets' to the Institutionalization of 'OpenScience' Stanford University Department of Economics Working papers 00-025 and 01-006. Online at http://www-econ.stanford.edu/faculty/workp/swp00025.pdf and http://www-econ.stanford.edu/faculty/workp/index.html.

De Soto, Hernando. 2000. *The Mystery of Capital*. Basic Books.

Dosi, Giovanni (1988), Sources, Procedures, and Microeconomic Effects of Innovation. *Journal of Economic Literature* 26:3 (Sept.), 1120-1171.

Foster, Richard N. 1986. Innovation: The Attacker's Advantage. McKinsey & Co.

Freiberger, Paul, and Michael Swaine (1984), *Fire in the Valley: the making of the personal computer*. McGraw-Hill, Inc.

Greenwood, Jeremy, and Mehmet Yorukoglu. (1997), "1974," *Carnegie-Rochester Series*, 1997. http://www.econ.rochester.edu/Faculty/GreenwoodPapers/1974.pdf

Harhoff, Dietmar, Joachim Henkel, and Eric von Hippel. (2002), "Profiting from voluntary spillovers: How users benefit by freely revealing their innovations." Working paper, May.

Jakab,.Peter.

Kirsch, David A. *The Electric Vehicle and the Burden of History.*

Lerner, Joshua, and Jean Tirole. (2002), "Some simple economics of open source." *Journal of Industrial Economics*, 52 (June).

Levy, Stephen (2001), *Hackers: Heroes of the Computer Revolution.* Penguin edition.

Liebeskind, Julia Porter, A. Oliver, Lynne Zucker, and Marilyn Brewer (1996),. "Social Networks, Learning, and Flexibility: Sourcing Scientific Knowledge in New Biotechnology Firms," *Organization Science*, 7:4, 428-443.

Meyer, Peter B. 2003 "Episodes of collective invention." U.S. Bureau of Labor Statistics Working paper WP-368. Online at http://www.bls.gov/ore/abstract/ec/ec030050.htm.

Mokyr, Joel, (1993), "Editor's Introduction," *The British Industrial Revolution: An Economic Perspective.* Westview Press.

Mokyr, Joel. *The Gifts of Athena*.

Nuvolari, Alessandro (2001), "Open Source Software Development: Some Historical Perspectives." ECIS working paper.

Nuvolari, Alessandro (2002), "Collective Invention during the British Industrial Revolution: the Case of the Cornish Pumping Engine." Forthcoming in *Cambridge Journal of Economics*.

Pavlicek, Russell C. (2000), *Embracing Insanity: Open Source Software Development*. Sams Publishing.

Raymond, Eric S. (2001), *The Cathedral & the Bazaar: Musings on linux and open source by an accidental revolutionary*. O'Reilly.

Saint-Paul, Gilles. (2003) "Information shring and cumulative innovation in business networks." CEPR Discussion Paper No. 4116

Schrader, Stephan (1991), "Informal technology transfer between firms: Cooperation through information trading," *Research Policy* 20: 153-170.

Shapiro and Varian. *Information Rules*.

Thurston, R.H. 1878. online, at Cornell.

Tushman, Michael L., and Philip Anderson (1986), "Technological Discontinuities and Organizational Environments." *Administrative Science Quarterly* 31 (Sept.): 439-465.

Usselman, Steven W. 2002. *Regulating Railroad Innovation: Business, technology, and politics in America, 1840-1920*. Cambridge University Press.

von Hippel, Eric (1987), "Cooperation between rivals: Informal know-how trading." *Research Policy* 16: 291-302.