

BLS WORKING PAPERS



U.S. DEPARTMENT OF LABOR
Bureau of Labor Statistics

OFFICE OF PRODUCTIVITY AND TECHNOLOGY

Episodes of Collective Invention

Peter B. Meyer, U.S. Bureau of Labor Statistics

Working Paper 368
August 2003

Results and conclusions in this exploratory research are the author's and are not official findings of the U.S. Bureau of Labor Statistics.

Episodes of collective invention

by Peter B. Meyer
meyer_p@bls.gov
U.S. Bureau of Labor Statistics
Office of Productivity and Technology
4 August 2003¹

Abstract:

The process of developing a new technology through open discussion has been called *collective invention*. This paper documents two episodes of collective invention and proposes a general model based on search theory. The first episode deals with the development of mass production steel in the U.S. (1866-1885), and the second with early personal computers (1975-1985). In both cases technical people openly discussed and sometimes shared technology they were developing. Both technologies advanced to the point that they supported substantial economic growth. Open source software development is partway through a similar process now.

The episodes have common features. The process begins with an invention or a change in legal restrictions. Hobbyists and startup firms experiment with practical methods of production and share their results through a social network. The members of the network form a new industry or change an existing one. The network then disappears if the new firms keep their research and development secret. A model of the search for innovations can describe this process if it is expanded to include independent hobbyists and consultants as well as profit-seeking firms.

¹ Thanks to Alessandro Nuvolari, John W. Meyer, Leo Sveikauskas, Thomas Maloney, Michael Giandrea, Joel Mokyr, Cindy Zoghi, Mike Harper, Marilyn Manser, Aklilu Zegeye, Anastasiya Osborne, and participants at the 2002 Social Science History Association meetings for their comments. Further comments are welcome; please address them to meyer_p@bls.gov.

1. Introduction

Technological advances are often kept secret or patented, making them the intellectual property of their inventors. Scientific advances are more often published openly. One reason for the difference is that scientific investigation is driven so much by curiosity, whereas technological investigation is clearly functional, driven by the goal of producing something and usually to earn a profit.

There are important exceptions. For example, the open, public character of open source software development² is unlike the privately-owned structures usually associated with technological development. This paper outlines historical episodes analogous to the open source phenomenon, and makes generalizations about them to build a theory of what happened. In particular it discusses the institutions that supported the open sharing of technological information in early U.S. steel production between 1867 and 1881, and those helping early microcomputer development from 1975 to 1980. These two examples are compared to three other cases which have been called *collective invention*. To think through a general theory it helps to have several cases of the phenomenon to be explained. With five cases, this paper outlines a generalized account of this phenomenon. The paper will make a useful contribution if it persuades readers that collective invention is a recurring phenomenon, that we know of several cases, and that it is especially likely and important when new technologies offer great opportunity.

Sharing can help move a technology forward into applications quickly because many developers benefit from one another's experiments. In each episode, the participating individuals have diverse motivations, but the cluster of reasons to experiment and share results is the same. In each case, key innovators were excited by the opportunity to do something new, were absorbed in technical challenges, hoped to do good for the world, hoped to become rich, and hoped to earn prestige and respect. Some of these descriptions are remarkably similar whether the individual was an early experimenter in mass production steel, steam engines, microcomputers, or open source software.³ The

² Software *source code* is the human-readable text which is translated automatically into a program of instructions the computer can run directly. Common source languages for this are C, C++, Java, Pascal, BASIC, and FORTRAN. *Free software development* or *open-source software development* is an approach in which the source code is made widely available on a computer network and its use, modification, and redistribution by others is welcome. Licenses vary; for examples see Pavlicek (2000) or <http://opensource.org>. The culture of open-source development is distinctive, differing importantly from corporate research and development, as discussed for example in Pavlicek (2000) and in the social science papers at <http://opensource.mit.edu>. Proprietary software producers typically hold the source code as a trade secret, and as private property, and distribute only an executable (runnable) form of the program.

³ There are many examples in the superb *Hackers* (Levy 1984), Freiburger and Swaine (1984), Cringely (1996), and McHugh (1980). Economists thinking formally about income and wealth as motivations, sometimes treat the experimenters behavior as idiosyncratic and surprising, but it is clear in the descriptions that the behavior is common, recurring, and can be understood by natural impulses to (a) do things that are interesting or fun, (b) earn prestige or respect, (c) try to get rich, and (d) do good for others. The subject of this paper is not why do they experiment but why do they publish their results? What problem does sharing help them solve? I argue that it looks like a search-for-innovations problem.

institutions through which the experimenters share their results differ among these cases. Such institutions include journals, clubs, consultancies, and job turnover.

The discussion below describes these institutions to show the common aspects among them and to develop a theory of the environments in which they appear.

1.1 The idea of collective invention

As societies have adapted to waves of new technology there have been phases where technological information was openly and enthusiastically shared. We will look here at certain cases where technological developments were published openly to some community, and we consider why it happened and what institutions made it work.

Robert C. Allen (1983, p. 2) used the term *collective invention* to describe “the free exchange of information about new techniques and plant designs among firms in an industry.” Allen saw that this had happened among iron producers in Britain’s Cleveland district. Richard Nelson (1982, p. 468) encouraged research into the subject. Schrader (1991) and von Hippel (1987) documented explicit, informal “know-how trading” among mini-mill steel makers in the U.S. in the 1980s. Nuvolari (2002) showed that steam engine engineers in Cornwall maintained a collective invention approach. Harhoff, Henkel, and von Hippel (2002) collected several recent examples of “freely revealed” innovations in recent years and modeled a game between innovators in which revealing could be an optimal choice by an innovator because it helps diffuse the technology. A number of investigators have described open-source software projects as examples of collective invention; relevant papers are at the web site opensource.mit.edu.

One might take the view that a technology is truly invented only once, and subsequent applications of it are innovations which are part of a separate process of diffusion. If so, collective invention is not the right term – collective innovation or user innovation would be better. But in the cases we will consider, the first versions of the technology are not applied for long. The invention as generally perceived afterward includes the many improvements from the early discussion. In these cases the form of the technology that turns out to matter is the one slowly invented over the course of years, so for this paper’s purposes we will say it was collectively invented.

There are several differences between collective invention environments and those depending on legal enforcement through patents. A patent system tolerates bogus submissions or unclearly described inventions, since an inventor benefits from users who pay for a license or are otherwise dependent. On the other hand, in a collective information environment, the innovator does not have an incentive to hide essential details. Patenting environments encourage the inventor to invest in and develop inventions, whereas collective information environments do not give an inventor a direct financial incentive to develop it further. Either environment could give innovators the satisfaction of seeing their innovations widely used. Which one will generate more

technological advance depends on context. Below I argue that when the technology's future is uncertain, collective invention does better at advancing the technology.

As we shall see, restricting the technical community to for-profit firms makes it hard to use the term in other episodes. Therefore let us enlarge the definition for the purposes of historical comparison:

- Include hobbyists and workers in not-for-profit organizations such as universities and government agencies in the community of technologists.
- Include job changes by people working with the technology as one of the mechanisms of interchange of technical information.
- Include cases in which collective invention practices co-exist with some firms who do secret research and development. Even if some firms keep their findings secret, collective invention may still occur in some subset of an industry.

So collective invention is defined here to be *a process in which improvements or experimental findings about a production process or tool are regularly shared*. Put this way, collective invention seems to be an important and regular feature of the historical process by which societies adapt to radically new technologies. It is part of a larger picture in which the new technology turns into new products and the producers, consumers, and markets are jointly developed.

Tushman and Anderson (1986) called the adjustment phase a “period of ferment,” describing the long period after an invention but before a substantial industry has stabilized around it. For example, the Bessemer steel-making process was first tried in the U.S. in 1858 but it took twenty years before there was there a strong industry built around it. Hobbyists made microcomputers as early as 1975, but it took until the 1990s before national productivity statistics were clearly affected. Open-source software by various names has existed since the 1960s but companies built around it appeared first in the 1990s. In each of these cases, the intervening period included an open, documented technical discussion. It could be that an open phase of this kind is necessary for adapting new technologies to useful purposes. Perhaps without the sharing phase, these technologies would not be developed to the point that they were beneficial to the general population.

2.0 Three established cases of collective invention

This section summarizes three large scale cases of collective invention documented by other authors.

2.1 Blast furnaces in Britain's Cleveland district

Allen (1983) found that from the 1850s through the 1870s iron-making companies in northeast England's Cleveland district allowed visitors and consultants to see the insides of plants and to write about the way their blast furnaces made usable iron from ore. Well-known researchers at the time such as Isaac Lowthian Bell, Thomas Whitwell, and J.G.

Beckton published information about the designs, size, temperature, and contents of blast furnaces. The information came from observing production, not from formal research efforts. Publications and well-informed consultants helped establish which blast furnace designs used fuel most efficiently. Plant designs evolved to have taller furnace stacks (filled with the input materials) and toward hotter and hotter furnaces. The design changes were not generally patentable because of the nature of the technology.

Since plant design was a natural area of competition one might have thought the owners of blast furnaces would each prefer to keep this information secret. But Allen concluded that through the sharing process firms could reasonably expect to learn more valuable information than they gave up, and therefore each firm preferred such information be made public over taking the risk of shutting it down by withdrawing. The collective invention regime substituted for research and development spending. Furnace efficiency improved over time. Allen found that little of this improvement was caused by private research and development efforts. He attributed most of the productivity improvement in blast furnace practice over time to the sharing of information which he called collective invention.

2.2 Steam engines discussed in Lean's Engine Reporter (1811-1904)

Steam engines had been around since 1712 but in 1769 James Watt patented a new, much more efficient design for them. Despite legal attacks on his patent, it was upheld until 1800. Mine owners in the Cornwall region of England used steam engines to pump water out of mines, sometimes using illegal copies of Watt's design. Mine owners and steam engine makers resented Watt's unwillingness to license the invention cheaply. After Watt's patent expired they could legally make modifications to the design. There was an explicit debate on alternative forms of intellectual property rights among the steam engine engineers. Few actually filed patents. Rather, there was a collective invention environment, as shown by Nuvolari (2002).

Starting in 1811 there was a publication read by the Cornish steam engine makers, called *Lean's Engine Reporter* for its editor, Joel Lean. Its contents were technical comparisons of operating steam engines. Nuvolari (2001) establishes that the efficiency of steam engines improved substantially in Cornwall through this period, probably through many minor or unattributed innovations and discoveries by the steam engine engineers. Collective invention sustained by the *Reporter* thus supported useful engineering improvements.

2.3 Open source software development since the 1980s

A contemporary example of collective invention is open source software development. In such a project the human-readable source code files are published on a computer network. Source code files, usually written in a standard computer language,

are given to specialized development tool programs, such as compilers, assemblers, interpreters, and linkers, which generate the machine-readable executable program.

Sharing the source code is useful insofar as it makes possible ongoing improvements by many programmers. Users may alter the program for their specific purposes. Sponsors of open source projects usually copyright the software in such a way that other developers cannot copyright programs using the open source code. This is a powerful mechanism to support collective invention because it is common knowledge that some later improvements will become part of the shared code.

The moderators of changes in a chunk of source code, also called its owners, determine the final choices in released versions of the software. Users may make a version different from the released, certified one. One criterion of a moderator's success is whether the moderator can avoid the project's source code "forking" into permanently divergent, partly-incompatible versions. If that happens, the mutual benefits of having one standard which improves over time are partly lost.

Several roles and institutions support collective invention in open source projects:

- Web servers are a venue for storing and distributing the technology.
- Special copyrights apply to most open-source projects, and intellectual property issues are confronted explicitly.⁴
- The relevant programmers have similar development tools and the skills needed to use them.
- Source control programs track which programmer is changing the software.
- Moderators decide and control which of those changes stay in the source code.
- Culturally, experimentation is welcome, and the developers are not scheduled or in other ways restricted from experimenting. This is one way that open source projects differ from many corporate projects.

Collective invention is easier in open-source software than in the earlier episodes. Software technology is quicker and easier to transfer over a network between participants than were the other technologies which were shared through documents. Also, shared source code is relatively easy to convert into a usable program, whereas a blueprint or description of a blast furnace cannot automatically generate a plant. In software, the content under discussion is the technology itself, not descriptions of it.

These three accounts, of blast furnaces, steam engines, and open source software, were drawn from an established literature. The next sections discuss two additional cases.

⁴ Pavlicek (2000) and opensource.org have examples of these special licenses to which a user must agree to legally use the source code. In the U.S. source code files are automatically copyrighted at the time they are first written, in the absence of specific rules otherwise.

3.0 Bessemer steel in the U.S. (1866-1885 and beyond)

British inventor Henry Bessemer announced a new steel making process in 1856. He correctly foresaw it would be a quick, high volume, fuel efficient approach and trod lightly on the fact that it hadn't actually worked yet. After some further innovations with others, it developed into a thriving business in the 1860s. It took longer to transplant the technology to the U.S., which took another wave of innovations.

Several institutions served to aid collective invention in the U.S. technology. First, several new industrial journals and organizations arose during this period when mass production steel technologies were being adopted. The open-discussion environment in the British iron and steel industry may have been a model for its extension in the U.S., leading to collective invention on both sides of the Atlantic. Second, famous engineer Alexander Holley ran a consulting practice which pooled the key patents for Bessemer steel manufacture so members of one licensing organization would have access to them all. Holley himself designed most of the first fifteen Bessemer steel plants though they had different owners. Third, job turnover was high in the industry, so many employees had diverse experiences drawn from previous employers.

Huge demand for steel rails for railroad construction sustained the industry through a depression that started in 1873. Production quantities rose dramatically, as shown in Figure 1. The technologies in use improved quickly, and the price of Bessemer steel fell from over \$100 per gross ton of rails in 1870 to about \$60 in 1880. Below we consider the institutions that enabled this collective invention.

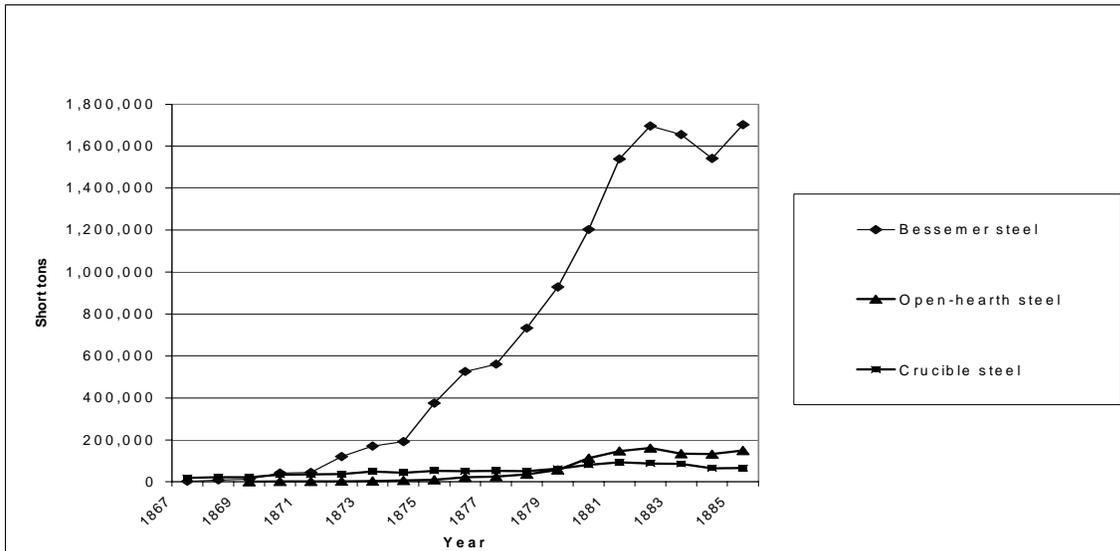


Figure 1. U.S. production of steel

Bessemer and acid open-hearth were new methods of steel production introduced in the late 1860s. Another method, called basic open hearth, was introduced in 1878. Steel production grew dramatically over the period. The data is from *Historical Statistics*, p. 694.

3.1 New journals and organizations

Metallurgical knowledge in U.S. industry was very imperfect in 1865. Standard textbooks on iron metallurgy in the U.S. had errors (Gordon, 1996, pp. 214-219, 297-8, and 314-316). It was not fully understood that the smelting process by which iron ore was converted to usable iron involved not only melting the iron but also chemical transformations which burned the carbon away. Iron and steel managers did not generally acknowledge the value of having chemists researching the subject, particularly in the U.S. By 1880, the situation had changed permanently. Standard textbooks had improved greatly and the subject of iron metallurgy was better divided into what was in fact known, and what was not understood but under study.

These improvements in knowledge in the industry occurred partly because of active investigation by professionals into iron and steel making, great demand from the railroads for iron products, and a set of iron and steel-making methods that was more diverse than before or afterward. Several professional associations were formed. Active discussion about metallurgy was documented in their journals. Institutionalizing innovation and communication in these ways could have helped change interaction to make it look more scientific rather than craft-oriented or industrial. Thus these journals and associations supported collective invention.

In Britain at the same time, the Iron and Steel Institute was formed and began to publish its periodical, the *Journal of the Iron and Steel Institute*, in 1869. It is possible that this publication inspired the American parallels, but it is more likely that the Americans were responding to a similar challenge, which was the need to understand and develop iron and steel technology better. Because of the new Bessemer, open hearth, and other processes there was more diversity in iron and steel production techniques than previously. The diversity later declined in a kind of shakeout.

The American Institute of Mining Engineers began in 1871 to publish the *Transactions of the American Institute of Mining Engineers* (TAIME). This journal was devoted principally to metallurgical subjects, especially iron and steel. Membership grew rapidly. The *Transactions* listed its members, about 700 in 1878 and about 1000 in 1880. To understand the nature of the publication, I reviewed the parts of 1878-1880 issues that discussed iron metallurgy and industrial processes.

The contents of the *Transactions* were technical, as opposed to discussing the affairs of the Institute, industrial personnel, or business matters. Readers were assumed to have a serious interest in the subject, and rarely did articles have an introduction or conclusion designed to draw in a passing reader. A number of the papers were presented at Institute conferences. Commentaries by discussants in the *Transactions* contained frank disagreements, politely but sharply written. The publication includes articles by directors of research, chemists, superintendents, and engineers. Membership included both PhDs (identified as such in the journal) and practical engineers who had not finished secondary

school.⁵ The text was formally written and readers were presumed to have some knowledge about the production and use of iron materials and rails, but not much of chemistry. Writers referred frequently to numerical evidence and to their experience. Many of the writers were well known figures in the industry rather than scientific specialists. The articles tended to be very empirical and specific about details of measurement.

Many articles had graphs or tables of data. Occasionally they had diagrams of plant or equipment. Some of the diagrams were so large that the reader had to unfold a page many times to see it completely. Other articles had a number of chemical formulas or mathematical equations. Diverse sources of specialized knowledge were welcome, clearly including input from steel users and customers.

Here are a few short examples of articles in the TAIME in the 1878-1880 period, and one extended example.

- There was a discussion of what technical training apprentices in technical professions should receive. Comparisons were made to the practices in Germany.
- A 94-page glossary of mining and metallurgical terms was published.
- The superintendent of a steel smelting department described specific improvements his firm had made to the ladles that held molten iron or steel for casting.⁶ A discussant identified one that seemed to be novel and important. Apparently it was not patented.
- On the basis of blast furnace evidence a professor showed evidence supporting a theory of the heat generated by combustion as a function of carbon and oxygen inputs and rejecting a widely held alternative theory.⁷
- Well-known engineer Alexander Holley described a new machine at the U.S. Arsenal in Watertown, Massachusetts which could test metal parts, and advocated that his fellow engineers try to persuade Congress to fund its availability to make machinery better and safer.⁸ His text took on some energy: “Machinery in vessels, on railways, in the floors of great factories and theatres plunge down among broken columns, torturing and killing men and women in their debris. Is it not probable that the tenth part of the money damages paid for these disasters, if expended in the means of prevention indicated . . . would very largely reduce this record of bankruptcy and death?” By its rules the AIME could not take a stance on a political question so Holley asked his colleagues to act as individuals and as representatives of companies.

⁵ For example, John Fritz, a noted inventor and the superintendent of one of the largest steel firms, had worked as a blacksmith’s apprentice starting at age sixteen. Engineer and inventor Captain Bill Jones published in the journal although he had apprenticed to an ironworks at the age of ten. These men seemed never to have returned to school (McHugh, p. 219 and p. 235).

⁶ Herrick, J.A. “Improvements in the appliances for venting molten steel or iron from a casting-ladle or shoe.” TAIME, 1878, pp. 13-15.

⁷ Church, John A. “The mode of combustion in the blast-furnace hearth.” TAIME, 1878, pp 33-44

⁸ “The United States testing machine at Watertown Arsenal”, TAIME, 1878, p. 261.

A kind of technical community had formed, with some warmth, like a freewheeling email discussion in the open source software context.

One particular set of articles is especially relevant to the developments in iron and steel. Pennsylvania Railroad's research director Charles B. Dudley wrote a lengthy paper on the subject of how impurities affected the survival of steel rails. The paper and comments by others were printed in the TAIME journal. The subject drew a dozen further commentaries in a later issue, and was central to the growing success of steelmakers in feeding the demands of the growing railroads. The superintendent of a prestigious and successful plant, Robert W. Hunt, wrote to say Dudley's sample was too small and to dispute specifics of Dudley's work with his own numerical support, and gently impugning the value of a laboratory contribution to an industry with practical experts such as himself. Well-known engineer W.R. Jones, of a competing plant, supported Dudley with certain reservations. Jones compared Dudley's findings to those of French metallurgists. He expressed the concern that American steelmakers were too secretive and shared what he thought was the best formula. A Columbia University professor deemphasized the importance of chemistry, which Dudley had studied, compared to the physical and mechanical microstructure of a rail in determining its robustness. This point of view was justified by later developments, according to Gordon (1996). Holley commended Dudley but recommended testing a sample a hundred times larger, and made a comparison to the practices of European rail makers. Another steel works engineer had another small sample of steel rails analyzed by three different chemists and submitted the results in his commentary. Another expert ran a least-squares regression, apparently by hand, on Dudley's data. There were 23 observations of three measures of reliability (the dependent variables) on the quantities of silicon, manganese, carbon, and phosphorus (the independent variables). A subsequent commentary criticized any analysis that did not take quantities of sulfur and copper into account. Several other professors and professionals wrote in with varying sources of authority. The discussion took over 130 pages.⁹

It is useful to note the diverse backgrounds of the participants, the diverse sources of information, and the process of this particular discussion. Participants included obscure experts as well as well-known ones. Attention was focused firmly on the problem of what made good steel railroads. The seriousness and enthusiasm of the participants suggests that they were making serious efforts to win the respect of one another through this process. A reader of this discussion can expect that the discussion would not get mired forever in detail. Through the open discussion, empirical findings would become established in a few years and a consensus on the facts or a design standard would emerge.

Other professional associations related to the developments in iron and steel appeared in the 1865-1880 period. In the late 1860s a professional association of civil engineers began. Alexander Holley co-founded the professional association of mechanical engineers in the 1870s. Its journal, the *Transactions of the American Society of*

⁹ "Discussion on steel rails," TAIME, 1879, pp 357-413 and 529-608. The series of articles reads like a thread on a Usenet bulletin board, but the phrasing is more formal.

Mechanical Engineers, began publication in 1879. And in 1880 the *Journal of the United States Association of Charcoal Iron Workers* began. Each of these had discussions about iron and steel materials, equipment, and processing.

3.2 The Bessemer patent pool and Holley's consulting practice

Litigation over the scope and definition of key patents held by inventors William Kelly, Henry Bessemer, and Robert F. Mushet delayed the development of Bessemer steel in the U.S. until after it was an established technology in Britain. A series of agreements beginning in 1866 defined a "patent pool" arrangement in which licenses to key U.S. patents were held in common by the Bessemer Association which paid fees through to the inventors and existing licensees. For \$5000, a firm could join the Association and not have to worry about specific patent rights. This was a small fraction of the setup costs of a plant (at least \$80,000) and was therefore not prohibitive. In practice member firms also licensed their own patents on improvements to the Association and thus its other members. American engineer Alexander Holley, a principal organizer of this patent pool, patented ten improvements in Bessemer plant design and process over the next fifteen years and made them available through the association. Holley published technical reports through the Association about approaches and experiments in applying these technologies. The presence of a newsletter did not mean it was easy to keep up, because much of essential knowledge was tacit, requiring experience with the new technology itself.

Holley's reports were intended for licensees only, but they may have been more widely available since the existing licensees would receive revenue from any new licensee. Holley was asked to join many projects. His role as an inspirational figure was analogous to some highly visible editors of early scientific journals. For example, entrepreneurial editor Frederick Oldenbourg personally financed the first major English scientific periodical *Philosophical Transactions* for many years starting in 1665 (Kronick, 1976). Holley also visited experts at their own plants, sometimes without warning, when he was thinking about a hard problem (McHugh, p. 230, note 20).

Holley's reports were brief and technical, and explicitly intended for his clients, not for widespread publication. One (Holley, 1875) described the steel works plant of the British firm Brown, Bayley, and Dixon. It began: "I will . . . confine my remarks to those features of machinery and practice which we may advantageously copy or avoid." The 18-page document was functional and instrumental, densely packed with descriptions of structures and equipment. Scores of measurements, percentages, weights, and costs for materials were listed. Holley thought highly of some aspects of the plant design: "The arrangement for removing slag, etc., from the pit is the best I have seen anywhere" (p. 4). Holley wrote that the pressure pump was troublesome, like other English pumps, and that with respect to the steel-making process: "There is no novelty whatever in the process as conducted here. Spectrum analysis is never used" (p. 12). It is not clear what Holley exchanged to the firm to get permission to see their plant, but perhaps he shared previous reports with them or provided consulting services.

In 1877, the eleven surviving Bessemer steelmaking firms in the U.S. combined patents to form a joint subsidiary, and ceased to advertise licenses for the remaining Bessemer patents.¹⁰ New entrants did eventually appear, but they had to pay higher royalties than the existing ones. Prices for steel and for rails continued to fall despite output quotas set by the rail pool because production costs fell and new sources of raw materials were discovered.

This is a case in which institutions of collective invention – Holley’s reports and consulting practice – were shut down by an increasingly oligopolistic industry. In 1881, Andrew Carnegie’s firm began acquisitions which resulted in the 1900 merger that created the giant firm U.S. Steel.

3.3 Job turnover

In the 1870s, iron and steel managers and technical experts often changed employers. Almost any experienced worker had worked somewhere else earlier (McHugh, 1980). Employers seem to have treated this as acceptable and usually remained on friendly terms with their former staff. Those who did not leave knew other engineers outside the firm. Indeed “the five or six top engineers of the industry [met frequently] to discuss common problems” (Temin, 1964, p. 133).

This is an aspect of collective invention even if the firms would have preferred that it did not happen. No one firm or person controls turnover. An ongoing flow of departures support a general collective invention regime, and this may be the result of institutional structures or an equilibrium outcome.

It has likewise been observed that Silicon Valley firms learn from one another through the rapid movement of employees and contractors between them, perhaps against their will. In social science research it has been found that employees with longer tenure at their current employer are more productive, but here we see a situation in which productivity-enhancing efforts to the region are aided by high turnover (thus *low* measured job tenure) for the individual at a particular job. This could be a general phenomenon in these environments or just arises from the selection of these perhaps ambitious individuals who seize new opportunities.

4.0 The Homebrew Computer Club (1975-1985)

When computer chips were first made small, powerful, and cheap enough to make it possible to build calculators and something like computers with them, a number of clubs of hobbyists appeared across the U.S. to do this and discuss it. The Homebrew Computer

¹⁰ Accused by a newspaper of forming a monopoly, industry titans pleaded (a) that they feared overproduction, (b) that they were being taken advantage of, and (c) that in any case they had not arrived at an agreement. It is clear that there was an attempt to agree, though its success is not known. All the facts of this paragraph are taken from Temin (1964), p. 175.

Club was an important one. It met at Stanford University starting in 1975, following the invention of the microprocessor but before there was a personal computer industry. Computer designer Lee Felsenstein moderated meetings. The group also had a newsletter. The Homebrew club was described vigorously by Freiburger and Swaine (1984) (henceforth FS84) and Levy (2001) from which the quotes below were drawn.

“The group had no official membership, no dues, and was open to everyone. The newsletter, offered free . . . became a pointer to information sources and a link between hobbyists.” (FS84, p. 106) Homebrew meetings included a presentation, often of a demonstration of a club member’s latest home creation. Then there was “the Random Access session, in which everyone scrambled around the auditorium to meet those they felt had interest in common with them. It worked brilliantly, and numerous companies were formed. A remarkable amount of information was exchanged at those meetings, and much information had to be exchanged; they were all in unfamiliar territory.” (FS84, p. 106) Members were drawn to the hands-on experience of making computers and understanding the component parts. Few focused on the theory of computing, or even the social effects of computing.

Hundreds attended regularly by the end of 1975.¹¹ Some members foresaw great potential for microcomputers. One member described himself and other Homebrewers as a pivotal in something like the industrial revolution.

The Homebrew Computer Club was not merely the spawning ground of many Silicon Valley microcomputer companies. It was also the intellectual nutrient in which they first swam. Presidents of competing companies and chief engineers would gather there to argue design philosophy and announce new products. Statements made at Homebrew changed the directions of corporations. Homebrew was a respected critic of microcomputer products. The Homebrewers were sharp, and could spot shoddy merchandise and items that were difficult to maintain. They blew the whistle on faulty equipment and meted out praise for solid engineering and convivial technologies. . . .

After Homebrew meetings, the most fanatical of the members went to a local beer-and-burger place known as the Oasis – everyone just called it ‘the O.’ They sat in wooden booths at wooden tables, surrounded by the deeply carved initials of generations, and drank beer and argued computer design. They ignored the fact that they were competitors. There were a lot of things to learn in developing this new kind of product, and they weren’t about to let economic issues get in the way of learning all they could. (FS84, pp. 108-111)

There were great differences in the talents, prior experiences, and resources of various members. Steve Wozniak’s early experience there was that he didn’t know enough: “The others were talking about the latest chips: the 8008 and the 8080. Woz felt lost. He hadn’t heard of them. But he had designed a video terminal and the club was interested. .

¹¹ Saxenian, p. 34. At one meeting it was estimated that there were 750 people attending. (FS84, p. 106). This author attended a Homebrew meeting in the late 1980s which had perhaps 150 attendees. By then most demonstrations were conducted by companies, not by hobbyists.

. . . 'It changed my life,' Woz would later recall. 'My interest was renewed, and every two weeks the club meeting was the big thing in my life.' . . . Woz couldn't afford an Altair [the first microcomputer kit], but he watched with fascination." (FS84, p. 211)

Later he built his own. Wozniak "brought his computer to Homebrew and passed out photocopies of his design so that others could duplicate it. Like a perfect hobbyist, Woz believed in sharing information. The other hobbyists were indeed impressed . . . He called his machine an Apple." (FS84, p. 212) Wozniak and Jobs formed the Apple Corporation and hired fellow Homebrew members Chris Espinosa and Randy Wigginton.

Writing about the innovative strength of Silicon Valley, Saxenian (1994, p. 34) identified the Homebrew Club as one of many expressions of a culture of sharing technology. That Silicon Valley culture had been described as free-flowing before the Homebrew Club existed and was not caused by the club. But the Homebrew Club was one important institution among many that supported collective invention. The products of this environment later supported tremendous economic growth. Not only did the club make technical information available and interpretable, but it was part of the general enterprise for innovators and potential innovators to meet. It was generally understood that the most effective way to benefit financially from an innovation was to get a product to market quickly, and to set engineering standards if possible. Legal protections like patents were viewed with suspicion and took more time to arrange and to use as defenses.

Homebrew evolved. Members who had started companies stopped coming, partly because keeping company secrets would be uncomfortable at Homebrew. Keeping secrets for private advantage violated what Levy (2001) called the Hacker Ethic -- that information should be freely available. From Levy (2001), p. 269:

No longer was it a struggle, a learning process, to make computers. So the pioneers of Homebrew, many of whom had switched from building computers to *manufacturing* computers, had not a common bond, but competition to maintain market share. It retarded Homebrew's time-honored practice of sharing all techniques, of refusing to recognize secrets, and of keeping information going in an unencumbered flow. . . . they had secrets to keep.

[One former Homebrewer said] ". . . people would ask you about the company, and you'd have to say, 'I can't tell you that.' I solved that the way other people did -- I didn't go. I didn't want to go and not tell people things. There would be no easy way out where you would feel good about that. . . .

It no longer was *essential* to go to meetings. Many of the people in companies like Apple, Processor Tech, and Cromemco were too damned busy. And the companies themselves provided the communities around which to share information. Apple was a good example. Steve Wozniak and his two young friends, Espinosa and Wigginton, were too busy with the young firm to keep going to Homebrew.

5.0 Common elements of these histories

Let us back out from the specifics of the steel case and the personal computers case, to think about the general phenomenon and how the term “collective invention” applies. In these cases, engineers were optimistic about the future of the technology but did not know for sure how it would develop. The future nature of products, production processes, and markets was not clear, or was not commonly known. Visions of that future varied and at least some were wrong. This is *technological uncertainty*, which characterizes the cases discussed above.¹²

To see this clearly, suppose instead that industry participants could forecast perfectly what technology they would be using ten years later. Then sharing technology through collective information-sharing institutions would be a bad idea -- sharing would not gain anything but could let trade secrets get out. Suppose a less-extreme thing, that industry participants didn't know the future course of the technology but were confident which research and development would find it. Then it would be most efficient and profitable for them to conduct that research and development privately, perhaps with partners, and to keep the results secret or to patent them. Collective invention does not seem to occur in contexts like those but rather when uncertainty is great, and the players cannot predict well at all.

These collective invention waves began with a new opportunity -- a new invention, or the expiration of a patent. Before the period of collective invention, either too little is known to interest those in search of a better product or process, or they are legally prevented from getting the resources to participate in the collective search. Once that opportunity is recognized, players are drawn in.

The collective search process creates an externality which creates a network effect: member A derives some indirect benefit and perhaps cost when member B joins. The expected gain of members probably grows with the size of the network. In the long run there could be gains to the larger society from the new technology through lower prices, employment opportunities, and the opportunity to tax new products.

A stylized time ordering of a wave of collective invention follows.

- A new opportunity opens from an invention, or the expiration of a patent.
- Then there may be a quiet period, in which some interested parties are aware of the uncertain potential of the newly adaptable invention.
- Collective invention institutions form, creating a social network.
- A flow of adaptations (microinventions) follows, from hobbyists and firms.
- New firms appear, seizing opportunities to apply the new technology for profit.

¹² A collective invention regime can exist in the absence of technological uncertainty. A firm might want to join a collective invention regime if it could not afford to conduct the required research, or if the payoff expected from conducting private research would not cover its costs. The information exchanges among mini-mill steel plants discussed by Schrader (1991) were of this type. But the cases discussed here have a different flavor.

- With the establishment of a profitable industry, technological uncertainty is reduced and the collective invention process evaporates. Surviving firms run private research and development. If that is expensive, few firms survive.

Consider the similarities across cases in Table 1.

Table 1. Episodes of collective invention

	Steam engine case	Cleveland district iron blast furnaces	U.S. mass production of steel case	Microcomputer club, Homebrew example	Open source case Linux example
Instigating or enabling events	Watt's patent expiration 1800		Patent pool agreement in U.S., 1867	Microprocessors available, 1971	Internet (circa 1970), AT&T breakup, (1984)
Common institutions or publications	Lean's Engine Reporter, 1811-1904	Books and consultants, 1850s-1870s	Prof. journals like TAIME and Bessemer Assoc. publications, starting 1871	Homebrew computer club meetings and newsletter, 1975-1980s	Internet bulletin boards and the Linux source code itself, starting 1991
Price of entry or restrictions on entry	Not available.	Informal inclusion	\$5000 till 1877, then \$80,000. (Temin)	Zero	Access to Internet bulletin boards
Tacit knowledge, skills, and prerequisite tools	Steam engine engineering or operation		Investment capital and background in furnaces or rolling mills	Practical electronics knowledge	Unix development software (compiler, linker)
Readers	Mine managers in Cornwall region of southwest England	Iron makers in Cleveland district of northeast England	Bessemer patent licensees in U.S.	Silicon Valley computer hobbyists	Unix programmers, connected by the Internet
# of contributors	Approximately two dozen		Dozens	Many dozens	Thousands (Pavlicek, p. 63)
# of readers			around 1000 AIME members	Several hundred	Thousands
Editor or moderator	Joel Lean, then his sons	Isaac Lowthian Bell, and others	AIME, other professional associations, and Holley	Lee Felsenstein, Gordon French	Linus Torvalds

In most of these cases there was a common publication all the participants could read, though they did not need to read it to be participants. Participants could not name all the other participants, as colleagues in a single laboratory could. Yet in these environments the networks seem to arise naturally, that is, not because of a powerful organizer or common prior interests, but because of interest in the subject.

Once technological and market uncertainty is in the air, some people believe they see an opportunity to do something new and better than ever before. People believe, sometimes correctly, that they have a unique contribution to make which can give them joy and perhaps wealth. Partly because of that sense of opportunity, collective invention

can arise naturally. Because there is no established market or technology, participants in a club or journal lose little by sharing what they know, but they could gain a lot from the contributions of others built onto their own designs or insights. So, some of them share freely. Participants search for improvements and share their findings.

Economic predictions are associated with technological uncertainty and therefore potentially with collective invention projects. Tushman and Anderson (1986), Dosi (1988), Rosenberg (1996), Greenwood and Yorukoglu (1997), and Meyer (2002) drew these links between technological uncertainty and its economic effects, using various terminologies:

- Technological uncertainty and perhaps collective invention coincide with heavy investment in research and development and therefore low profits in the industry.
- Income inequality in affected sectors rises temporarily because opportunities for financial success and failure are created by the uncertain situation. Analogously on the profit side, many firms are wiped out even as others grow rapidly.
- Productivity improvement, as measured by current inputs and output, slows down as the industry or economy adapts to the new situation. Productivity statistics undercount useful work done by collective invention since information is generated and given at a low cost. But measured productivity may rise greatly after engineering standards are set and the relevant industries and technologies are established.

These economic observations generally do apply to the steel case and the microcomputer case discussed above.

5.1 The social network perspective

Open-source software projects can be modeled as social networks, made up of developers with links to the project or to the other developers.¹³ To stay consistent with the theoretical literatures, members of social network are called *actors*, whereas in search contexts they are called *agents* or, in a game theory context, *players*. These terms always describe the technology's developers. The earlier cases of steel development and microcomputer creation could also be described this way.

Liebeskind et al (1996) made a useful definition of social networks applicable to information interchanges like those in collective invention environments. "A social

¹³ Examples include Cowan and Jonard (2000), Madey, Freeh, Tynan, and Hoffman (2003), and the sources cited in Scacchi (2003, p. 19). An empirical regularity of such networks shown by Madey, Freeh, and Tynan (2002), is that power laws describe the distributions of (a) the number of projects a particular developer joins, and (b) the number of projects with n developers. If one thinks of each previous historical episode as one case, the number of examples given in this paper is too small to test this hypothesis on the cases presented. A more precise approach would be to take each journal, in the steel case, or each club in the microcomputer case, and treat it as a network. That would also reduce the serious selectivity problem in this paper, which discusses only those cases of collective invention which succeeded on a large scale and were well known after the fact.

network can be defined as *a collectivity of individuals among whom exchanges take place that are supported only by shared norms of trustworthy behavior.*” Communication through networks is thus more like scientific communication (exploratory, and not explicitly paid) than like commercial communication. They went on to compare networks to hierarchies and to markets compactly and elegantly. Three components are relevant here:

- “Unlike hierarchies, but like markets, social networks involve exchanges between legally distinct entities.” Network links are external to each organization.
- “Unlike markets, but like hierarchies, social networks support exchanges without using competitive pricing or legal contracting.” They depend instead on shared norms.
- “Social networks can [enable organizations to flexibly] switch from one source of knowledge to another without incurring the costs or commitments inherent in either hierarchical or market exchanges.”¹⁴

In a collective invention network, actors contribute streams of information into pools which may then be controlled by editors or moderators. Members can include hobbyists, employees, or corporations willing to share. For sharing to be useful the participants must have different expertise, experiences, or tools. Useful knowledge in a collective invention network accumulates over time.

Actors may share their findings because that is ethical behavior, and they feel good about it. They also build a reputation, which provides psychic benefits, prestige and entrepreneurial opportunity. The cost of sharing is sometimes low, especially in the open-source case. In the open-source case each actor can help ensure that the benefits of future fixes and features developed by others apply to the player’s own system by giving the others enough code so they are running systems compatible with the player’s own. Lurkers (non-contributors) impose few costs and may provide future benefits.

One kind of network is an *information brokerage*, which is a process “by which intermediary actors facilitate transactions between other actors lacking access to or trust in one another” (Marsden, 1982). Arrangements other than a star-shaped network, in which an actor at the middle of the network controls what new information is distributed to the others, would delay information transmission, and introduce new intermediaries and associated principal-agent problems. (Gould and Fernandez, 1989). The editor of the journal, at the center of the star-shape, is serving many authors and readers. This induces principal-agent problems, because the editor is an agent of so many principals, but this person is monitored by the entire population of players. So the arrangement with one monitored central editor publishing a journal is perhaps the most efficient shape of the network, or may minimize principal-agent problems relative to some other structure. It may be acceptable to the other players that the central player gets rich; for them the key thing is to stay informed.

¹⁴ Liebeskind, Oliver, Zucker, and Brewer (1996) pp. 430-439. Their subject was the development of biotechnology across universities and firms. This was a collective invention process too, with dramatic technological uncertainty, but they did not use these terms.

In both the network of steel developers and the network of computer builders, certain individuals were central. Alexander Holley was a charismatic central figure of steel development who wrote frequently in the journals. He encouraged participation, suggested appropriate behavior, and actively traveled and discussed the technology. He was an enabler for their work. Lee Felsenstein was the principal moderator of the Homebrew club, and made sure people not only listened to presentations but also talked to one another. A related advantage of the collective invention process in principle is that an extremely insightful or productive specialist is made useful to all the developers. In a privatized R&D structure only one cluster of developers would get that person's feedback.

Actors transmit messages containing partial results from experiments through the network and privately. They receive messages that pass through the filter provided by the newsletter editor, or private communications specifically addressed to one another. There is a cost to sending and receiving messages, but within the network it is low.

The information brokerage process with the star shape is not characteristic of diffusion, in which a long, slow sequence of adoption interactions take place. Here, the activity under study is not the diffusion of any particular innovation, but the process of examining one after another. The information brokerage structure does not strictly describe the innovation process because it is clear in each historical case that important conversations between happen outside the newsletter too.

It does not follow from this process that the best ideas prevail, although they do have an advantage. A design could dominate even if it were not the most efficient. The choice may be made early (as in the QWERTY typewriter keyboard case) and not be reconsidered since the members of the network are not organized in such a way as to make a collective decision.

The information brokerage structure is not competitive, so as long as it is oriented around a hobbyist research problem it is not properly modeled as an industry. But as the technology improves, there could be profits to be earned from the technology, and then the interests of members become differentiated. Some clubs may become firms. That is a club could incorporate and commit itself to using some particular technology and equipment, in order to make profits. Firms or other subgroups could decide to keep their information secret. Without contributions from these subgroups, other players may lose interest in the network. So if members break away, the network may collapse from lack of interest or resources. If the member firms competed for profits the network could break up and become a competitive industry. The network could thus start an industry whether the network continues on or not. Homebrew alumni developed several personal computer standards including the Apple II, the Osborne, and the Macintosh, which were competing products. The Homebrew club survived into the 1990s and the Apple Corporation continues to sell descendants of the original Macintosh.

In the cases we have considered, important innovations come from people within the network, which indeed is best understood as a network not a market or hierarchical

phenomenon in the language of Liesbeskind et al. A model can leave out people in the economy who are not in the network. But studying the network per se leaves out some deeper questions. Let us think out next why the individuals want to create the network.

6.0 The search for innovations

The network exists, I argue, because it reduces the costs of searching. The institution is a kind of search equilibrium outcome. Participants build the relationships in the network because of particular conditions in the environment. We can describe these in supply and demand terms. In brief, participants demand flows of information to improve their technology because they believe there is an opportunity to do something new and useful. There is a supply of information about relevant innovations because the technology is in fact new and people are trying to figure it out. Both of these result from technological novelty and uncertainty.

A market could appear for this information. This is more likely if the information is immediately useful in a predictable way in production, particularly to profit-seeking firms. It is also more likely if the bits of information are homogenous and therefore can have a standard price. But under conditions of great technological uncertainty these conditions do not apply. We observe that many of the participants are not directly profit-seeking, and that the bits of information going through the network are heterogenous and unique, and that they are not priced. We discuss how to model these attributes below.

6.1 Seekers and the innovative search

Search theory in economics normally describes agents who sample a distribution repeatedly, stopping optimally when further search appears on average unprofitable. The search might be for job opportunities, applicants for a job, or marriage partners. Search theory has been used to algebraically characterize searches for innovations too. In search models of innovation, a profit-seeking enterprise decides when to conduct formal research and development (R&D). A sophisticated example is in Jovanovic and Rob (1990). In these models innovative search is a gamble whose possible payoffs are averaged together, and if the average outcome is profitable, the firm devotes resources to the attempt to innovate. The firm weighs the costs of an experiment or search attempt and the distribution of possible profits from it. The balance or equilibrium between these in the model predicts an amount of R&D effort that is more than zero but not infinite.

What would have to change to make such theories applicable to the early periods of innovative search? Several possible changes to these R&D models could adapt them to describing the early search for a useful form of a new technology.

6.2 The subsidy implicit in enthusiasm

Allen (1983) and others have offered many hypotheses to explain why self-interested private firms or individuals would release proprietary information. The list can be overwhelming. Allen (1983) has examples supporting the first six.

- First, hobbyists and people within firms have ambitions that they can help meet by releasing technical information – prestige, fame, or employment opportunities.
- Second, firms garner publicity by making their successes known, and it motivates their staff to compete to improve.
- Third, it can be costly to keep information secret especially when there is substantial movement of employees between firms.
- Fourth, a firm’s release of information could increase the value of some asset it had. For example when better ironmaking methods were introduced, ore deposits owned by a British iron firm gained value, whether or not that firm used the new methods.
- A related (fifth) idea is that cooperating with other firms to improve production might be expected to induce some improvement to one’s own firm, 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.
- 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 sort of certification signal as suggested by P. David (1998).

Hobbyists have other reasons, too:

- For them, playing with the new technology can be fun and absorbing, and many hobbyists thought it was virtuous¹⁵ and exciting to share their findings.
- Lastly, 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.¹⁶

Combined, these motivations over-explain the phenomenon. There are many reasons to share information, and a state of technological uncertainty and opportunity contributes to most of them. In a static industry with unchanging technology in which the players anticipated little change, these incentives would not apply. But once the opportunity

¹⁵ 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.

¹⁶ This is Eric Raymond’s vivid phrasing, from <http://catb.org/~esr/writings/cathedral-bazaar/cathedral-bazaar/ar01s02.html>, and later in Raymond (2001), p. 23.

appears, there are many reasons to participate in sharing mechanisms for their information. Sharing and experimenting are complementary. Experimenters would be more likely to give up one if they did not have the opportunity to do the other.

In the early research discussed above, there were just a few individual experimenters. Most of the early experimenters lose money. In modeling terms this is different from the R&D equilibrium of an innovation search theory. Experimenters subsidize the search with their own enthusiasm and other resources. The enthusiasm flows from the belief that there is a great opportunity. This subsidy does not prevent a search model from making predictions. For example, as in the other models, there would be more searching if costs of experimenting were lower. It is possible the model could predict that less searching would be subsidized as profits appear and profit-minded R&D explores the unknown possibilities, displacing the hobbyists.

The subsidy can be measured sometimes. Moskowitz and Vissing-Jørgensen (2001) found that financial returns to entrepreneurs they studied were lower on average than a risk-reward tradeoff with stock markets would imply, once one takes into account that entrepreneurs were so often poorly diversified, with most of their holdings in their own firm, and were therefore carrying extra risk. These are observations of a kind of subsidy, made perhaps because the entrepreneur is focused on other things, or is very optimistic about the future of a particular venture, or is institutionally unable to diversify from it.

Consider all the factors that historically have motivated researchers into early technologies that are exogenous to the search discussion. Some, like Bessemer, underestimate the difficulty of turning the technology to a profit. Others are driven by prestige, not the desire to enhance production or earn a profit. Some believe they can change the world, and see this as an opportunity, rather than the opportunity to start a firm. Tim Berners-Lee, the inventor of the Web, seems to have taken this perspective. Like some other technology developers, Berners-Lee also had a particular application – sharing physics research among government institutions – for which it was reasonable to think the Web would pay off straightforwardly. Some early innovators see that there is an opportunity to set engineering standards or to be one of the first to sell a product, and purposely forego any efforts to obtain intellectual property protection. A number, in the Homebrew case, understood that once they developed relevant skills and were recognized in the new arena of microcomputer hardware and software, they would be very employable afterward regardless of whether their own design or innovation were valuable. This list is not complete – in fact, the motivations of early technology developers are many and diverse. So for purpose of modeling a search equilibrium in which the players share their findings, the experiments are exogenously subsidized.

6.3 Experiments that can create productive capital

Members of the social network may have weak links to one another which support potentially valuable matches. Granovetter (1973) found in a case study that a wide network of friends seemed to be more effective in finding a job than a few close, committed friends. Contacts in the network can help members reach desired goals.

Members discuss problems and possible solutions with one another. In a matching process, members evaluate whether others have useful tools and skills. The interactions have other dimensions as well, but these are key components.

Members of the network conduct private experiments based on their own resources, interests, and insights. This is productive work but does not generally increase current output because it is experimental.¹⁷ It can have financial value immediately however and eventually improve productivity. See Appendix A for an example of how to treat the experiments as productive, in expectation.

A player might have nothing to share before experimenting. The player might not experiment if there were not a venue in which to share the findings. Sharing and experimenting are complementary activities to achieve the individual's objectives. Each person may have some objective other than to search. For example, a player may want an operating system which gives the user control over which processes running on the computer have high priority, and the player is willing to edit the operating system code to achieve it. This person might not wish to have to make a better operating system, but may be willing to go that direction if that is the only way to get the better operating system. There is a collective search process, and *in the search, sharing and experimentation are complementary inputs*. Each person who wants to make progress in the search, experiments more and finds that it is optimal to share it. And players who participate in sharing more, find it is optimal to experiment more.

In the open-source software context, Lerner and Tirole (2002) asked “Why should thousands of top-notch programmers contribute freely to the provision of a public good?” There is an answer, here. The interpretation here is that programmers who do this would have experimented with the technology anyway, and they share because this makes experimentation more efficient. As posed by Lerner and Tirole, the situation sounds like one in which the programmer is giving capital or labor away. If we assume that the programmer has some objective to be met by experimenting with the technology, the programmer's decision to share seems more natural, and in all these historical cases the ones who share were experimenting. Within the context of all the experimental searches for innovations, there will be programmers, hackers or other technology developers who follow a “freely revealing” strategy (Harhoff, Henkel, and von Hippel, 2002).

6.4 A production function in which pure innovation is possible

Let $f(X, T)$ be a production function. Its value is the amount of some kind of output – in the Bessemer case an amount of usable steel. X is a vector of current input material quantities, including pig iron, iron ore, and the manganese compound called spiegeleisen which turned out to be an essential ingredient. T is a vector of measures of design and

¹⁷ An individual's purposes need not match the phrasing here for this description to fit. What seem like experiments in this structure may have entirely predictable consequences to the actor. An insight may be erroneous, but lead to a useful experiment. The technology may seem usable to some experimenters but not others. And many dimensions of X and T may not be clear to anyone at the time of an experiment. The history may be chaotic, but the description here imposes an after-the-fact perspective and coherence.

technological attributes such as the height of the furnace, and the duration of the air blow through it. T is held constant in the classic economic problem of optimizing output quantity for each set of inputs, or for minimizing inputs to obtain a fixed quantity of output, or for minimizing input costs to obtain a certain level of outputs. Let us assume here that optimizing over X is not too difficult in principle (whether or not it is feasible with the resources available). Instead let us concentrate on the problem of searching the space of possible T s. For example, the agents might seek the steelmaking design which produces the best steel output given as input one ton of pig iron.

The search for innovative adaptations of a new production process is like a search for a desirable T in a region of R^N , though the searcher may not know the number of dimensions N . Each point in the region represents a particular T and therefore a production function. Each dimension characterizes the choices for some attribute of the production function. Bessemer converters made steel by blowing air through molten pig iron so such design dimensions might include the height of the converter or furnace, the quantity of manganese added (which was zero by default until it was discovered that it improved the product), the duration of the air blast, and whether and how to integrate a blast furnace (to make molten iron) into a steel production process.

Let us assume members share an optimistic belief that there are more useful or profitable locations of $f(X,T)$ than have yet been seen. Each player has a private image of $f(X, T)$, based on imperfect information, which we might call $f_{i,t}(X,T)$ to index it by player and time. Most of the information is missing and some is untested or could be incorrect. A particular hobbyist-agent may have some goal other than learning about $f(X,T)$ but read the journal anyway.¹⁸ The presence of a journal about $f(X,T)$ focuses attention on specific technological problems, and makes it feasible to share news about $f(X,T)$. The resulting social network of readers and authors is not an organism with its own objectives, but an institution that helps hobbyist-agents achieve their own ends.

Historians of technology treat “success” and “failure” as ambiguous and socially constructed terms. (Lipartito, 2003, pp 53-58) A technology can be a failure in the technical sense that it didn’t work; in the financial sense that it did not pay off to its inventor; or in the market share sense that it was dominated by some other technology. Along with these three dimensions of success, there is another here. A technological experiment exposes the behavior of the technology with the design choices in the particular vectors X and T that the experimenter attempted. If there is a social network through which the result is shared with others, the experiment can help the group find success. Thus the ambiguity of success and failure here in this model of technological history is consistent with the ambiguity in the world the model should describe.

The seekers in the network learning about $f(X,T)$ are members of a larger population. Because the technology is uncertain, some people do not believe that $f(X,T)$ can be

¹⁸ Each player has private information of various kinds. Each player has psychological and financial objectives which might be met by combining private and public information. A field of science can be viewed this way; Joel Mokyr suggested the term “collective discovery” to parallel “collective invention.”

improved, but do not affect the history of the search. Indeed they could be forgotten by history.¹⁹ Many of the searchers are forgotten too.

In some regions $f(X, T)$ is smooth and single-peaked, so analytic optimization over T is possible if the function can be estimated. But $f(X, T)$ is not defined at every point, and has discontinuities. For discrete choices, $f()$ is defined only for certain T values, e.g., 1 or 0 for yes or no. An experiment may be a shot in the dark which lights up little space.

At the beginning not much is known about the dimensions of the vector T . This induces technological uncertainty among the network members and between themselves and the larger population. Uncertainty is intrinsic to the situation since the design problem is hard and not well known. In Table 2, the issues defining the dimensions of T were each uncertain in the empirical sense that at some time some participants bet on different choices but eventually one choice became an industry standard.

Table 2. Searches for technological and organizational design

	Bessemer steel, 1856-1881	microcomputer hardware, 1971-1982	open-source operating system
Output y, where $y=f(X,T)$	Quality-adjusted quantity of usable steel	Quality-adjusted capability of computer hardware to store, compute, and communicate	Ease-of-use- and quality-adjusted capabilities
dimensions of X (inputs to production)	Rates of input of ingredients: iron, speigeleisen;	Is it made from standard electronic parts? Were they new or used? Does it have an operating system?	Can it multitask? How well can it share files with other operating systems?
dimensions of T (design attributes of production)	Is manganese added? (0/1) How near is the steel plant to the Great Lakes by water? How near is it to low-phosphorus iron ore? How near is it to coal? Did a chemist work there? Is iron of several types mixed together? Is steelmaking integrated with iron-making in blast furnaces?	Is the architecture open? Can hackers make compatible hardware? Is it a smart terminal, or does it stand alone? Were programs stored on paper tape, magnetic tape, or disks? Can a modem plug in? Do customers want a smart terminal, or a standalone computer? Are graphics possible? Are screen pixels mapped directly to memory locations?	Which programmers can work on it? Do they use source control? Can a company make money selling software, or by offering services? Which licensing scheme is used? How does the moderator prevent the forking of one project into two?

¹⁹ The histories of collective invention that are best remembered are those which succeeded in a big way, such as the examples of steel and the personal computer. Ideally a theory here would have a way of describing the process when the invention did not succeed, or when there were people who thought it would not. A good example is that of the hot-air balloon, as discussed in Mokyr (1990). Some people thought this dramatic new invention would be useful and the technology was explored in France in the late 1800s. But no greatly profitable application was invented.

The definition of $f(X, T)$ is in terms of input measures of various kinds and of output quantities. It is not known at the beginning of the search whether $f(X, T)$ could be profitable. Different views may exist on which aspects of an existing technology are problematic. After the fact, a historian may confidently write that the problem was that the iron was excessively oxidized and the solution was to add a manganese compound, but beforehand the definition of the problem might well be subject to dispute. “Problems” are defined in discussions which include visions of possible “solutions.” Once a better design is clear, it becomes common knowledge that there was a problem which was solved. Until then, the process is not definitively known to be improvable.

The players have diverse resources, opportunities, insights, abilities, interests, skills, and agendas. Each one may have something unique to bring to an experimental opportunity, so pairs of participants can do experiments that the individual participants could not do. They have different capabilities in this exploration production function which produces stochastic results. If they hit a result that pays off, they may spin out into a business, and search $f(X, T)$ further on their own, keeping their own experimental results secret, or patent them. Capabilities that individuals develop from experience may outlast the experiments too.

The complete $f(X, T)$ is never known, because it is too complicated, many aspects are of no interest, and disputes remain. Other problems take over, such as profit maximization, satisficing for the survival of an organization or an employee. So the history we see does not include an explicit declaration of $f(X, T)$ whose mysteries are all uncovered; rather, we see a series of technologies and experiments. Errors are corrected, anomalous results investigated, and doubtful results established to general satisfaction. Players have ways of extrapolating information and beliefs about $f(X, T)$ into the space that is not known. Clubs may form around various visions of the future discoveries about $f(X, T)$. A club may then become a business.

Once a number of clubs form, the network may collapse or change its focus. The project of exploring $f(X, T)$ may no longer be of interest to a wide readership since (a) the best information about $f(X, T)$ is now private property; (b) there is little prospect of new entry against established competition; (c) there is no purpose in being a hobbyist any more since progress in the industry’s technology is sustained by profit. Technological uncertainty has then been reduced enough that the hobbyists are gone and search over T has been taken over by corporate R&D.

6.5 The option of creating new firms

As increasingly profitable locations in the domain of $f(X, T)$ are found, corporate R&D is more likely to appear. The resources of a profitable producer make it possible to build a patent portfolio and a hierarchical organization which is private property. Such an organization explores industrial dimensions of production outside $f()$, which describes production technological choices only. Among the organization’s choices are: to define the product; how to sell it; how to describe it in advertising; what customers to seek; what prices to charge; whether to behave monopolistically or litigiously; whether to patent;

whether to keep secrets, and so forth. A firm making such choices often does so privately, and if several do this, the collective invention network could evaporate.

There are examples of the collective invention regime disappearing. In the example of U.S. Bessemer steel development, the industry locked down its patents and charged a high price to new entrants. Large firms merged starting in 1881 and for the rest of the century. Around 1900 the last giant merger formed the monolithic U.S. Steel company. In the microcomputer case, many companies arose out of the Homebrew Computer Club, including Apple and Cromemco. These companies had secret projects, such as the one at Apple that produced Macintosh computers, and Homebrew lost its central importance. In the steam engine case, the developments shared in *Lean's Engine Reporter* eventually slowed. The *Reporter* ended publication in 1904.

6.6 Prospects for innovation search theory

The modeling approaches described above (hobbyists subsidizing the search; a production function with design and technological attributes; experiments that create capital; profitable departures from the network) could lead to an expanded version of innovation search theory in which (for example) restrictions on free speech, publication, or reductions in the scale of market production would reduce collective invention. Expansions of patentability could also reduce it. Technologies like the Web that make information-sharing easier should increase effort devoted to collective invention in a model. In such a theory time alone could make the social network collapse because its participants could break out to run profitable, secretive, enterprises. As described here, collective invention periods are valuable to society at large but are delicate, because they depend on the subsidy of the enthusiasts. The institutions are most important when there is uncertainty along many dimensions, which creates the possibility of truly major innovations.

7.0 Conclusion

Collective invention institutions discussed here include clubs, informal newsletters, formal journals, conferences, consultancies, anti-copyright rules, patent pools, engineering standards, and whatever institutions sustain high job turnover in technologically uncertain times. These institutions sustained social networks through which findings and innovations were shared without formal intellectual property rights to restrict them.

The incomplete framework here defines an innovative search problem and a social network of experimenters who try new production processes. Sharing and experimentation are complementary inputs to the search process. The network or sharing process is a way of searching more efficiently. The participants choose to search for a variety of reasons. Some believe they can get rich, and a few do. The network could break up into an industry of profit-seeking enterprises.

Although hobbyist activity can be productive, government actions rarely measure it or support it. Government policy can however help or impede collective invention processes. For example, governments have a variety of ways to encourage or prevent the interchange of electronic information, movement between jobs, and new company startups.

Productivity statistics do not count informal or unpaid research. In fact, if a hobbyist invests time and material inputs into an experiment, the hobbyist's costs could be counted as consumption even if the resulting new computer started an industry. Hobbyist, collective, or open-source efforts may be an essential phase of technology improvement. Early automobiles and airplanes seem to have developed along a collective invention path. Perhaps all capital goods industries experience this phase and it could be detected in productivity statistics.

If the institutions in the environment did not support a phase of collective invention, some technological developments would stall, and some technologies would never develop to become generally useful. Put this way, the idea is obvious. But this tells us something about where in today's world technological developments actually occur. It also helps us understand where technological development occurred historically. Technical journals, a free press, hobbyists, and collective enthusiasm are hardly mentioned in the classic historical works comparing the economic development of the medieval west to medieval Islam or dynastic China. The collective invention idea is not a major component of these accounts, but perhaps it needs to be.

This paper contributes to the literature on collective invention in several ways. It documents two more cases of the phenomenon – early steelmaking and personal computer building. It systematically lists similarities between all the cases, and it proposes a search theory approach to relate these similarities to the environments that produce them. The idea that experimenters personally subsidize the search is new and may be an essential component of models that explain the “period of ferment” in which a radical new technology finds practical applications.

Appendix A.

The financial value of an experiment can be modeled this way. Suppose a piece of equipment produces q units of output every period, the price of a unit of output is fixed at one, and the interest rate at which an employer can borrow or lend is fixed at r (e.g., .05 per year). Define the present value of a worker's output to be V on the basis of forecasted prices of inputs and outputs. The relation between the present value of the equipment and its output each period can then be shown as an annuity:

$$V = q + \frac{q}{1+r} + \frac{q}{(1+r)^2} + \frac{q}{(1+r)^3} + \dots$$

The terms after the first are equal to $(1+r)^{-1}V$ so a more compact equation is possible:

$$V - \left(\frac{1}{1+r} \right) V = q \quad \Rightarrow \quad V = \frac{q(1+r)}{r}$$

Consider now the value of having conducted an experiment that discovered a design change which would raise output. Let us say it costs c to implement the change, takes time t to have an effect on production, and raises output by Δq . The present value of this experiment is

$$V_{\text{improvement}} = \frac{\Delta q(1+r)}{r} \frac{1}{(1+r)^t} - c = \frac{\Delta q}{r} \frac{1}{(1+r)^{t-1}} - c$$

Estimates of c , Δq , and t at the time of the experiment may naturally be poor.

If the hobbyist pays (through the subsidy engendered by enthusiasm) a cost $c_{\text{experiment}}$ to conduct the experiment, and its chance of producing the output improvement is $p_{\text{improvement}}$, then the mean value of the experiment to a producer of the good, in advance of knowing whether it was successful, is

$$V_{\text{experiment}} = p_{\text{improvement}} V_{\text{improvement}} - c_{\text{experiment}}$$

There is also a social gain not modeled here. Experiments help the collective search by expanding the knowledge of T-space. Both successful and unsuccessful experiments therefore raise the probability that future experiments will succeed.²⁰ This value is increased much more if the information is publicly disclosed to other possible experimenters than if it is secret.

The value of experiments is greatest when there has been a potentially great new invention, and technological uncertainty is in the air – that is, the possibility is real that better locations within T-space exist than have yet been found. It has been argued that periods like this are associated with rises in income inequality and asset return inequality and volatility, because both predicted and after-the-fact values of experimenting vary so much.²¹ Holding all else constant, the value of experimenting is lower if r and t are high. High interest rates discourage experimenting, and so do institutions that would delay implementation of a technological improvement, e.g. some kinds of regulation or collective bargaining.

²⁰ Nelson (1982, pp 462 and 466) discusses a class of models in with this property, in the context of profit-minded firms. There, “knowledge means ability to guide R&D effectively” by advancing more for a given cost or reducing costs for a fixed advance. More generally, different kinds of knowledge “all provide clues as to where to search next.”

²¹ Greenwood and Yorugoklu (1997) and Meyer (2002) are examples of the dispersed literature on this topic.

References / Bibliography

- Allen, Robert C. "Collective invention." *Journal of Economic Behavior and Organization* 4 (1983): 1-24.
- Cowan, Robin, and Nicolas Jonard. 2000. "The Dynamics of Collective Invention." MERIT Working Paper 2000-018. Forthcoming in *Journal of Economic Behaviour and Organization*.
- 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.)
- David, Paul A. 1998. "Common Agency Contracting and the Emergence of 'Open Science' Institutions." *American Economic Review*, 88:2 (May), pp. 15-21.
- Dosi, Giovanni. 1988. Sources, Procedures, and Microeconomic Effects of Innovation. *Journal of Economic Literature* 26:3 (Sept.), 1120-1171.
- Freiberger, Paul, and Michael Swaine. 1984. *Fire in the Valley: the making of the personal computer*. McGraw-Hill, Inc.
- Gordon, Robert B. 1996. *American Iron, 1607-1900*. Johns Hopkins University Press.
- Greenwood, Jeremy, and Mehmet Yorukoglu. "1974," *Carnegie-Rochester Series*, 1997. <http://www.econ.rochester.edu/Faculty/GreenwoodPapers/1974.pdf>
- Gould, Roger V., and Roberto M. Fernandez. "Structures of Mediation: A Formal Approach to Brokerage in Transaction Networks," *Sociological Methodology*, 19 (1989): 89-126.
- Granovetter, Mark. 1973. "The strength of weak ties." *American Journal of Sociology*, 78:1360-1380.
- Harhoff, Dietmar, Joachim Henkel, and Eric von Hippel. "Profiting from voluntary spillovers: How users benefit by freely revealing their innovations." May 2002 working paper.
- Historical Statistics of the United States: Colonial times to 1970*. 1975. Bureau of the Census, Washington, DC.
- Holley, Alexander. 1875. *Brown, Bayley, and Dixon*. Printed privately. Available at the U.S. Library of Congress.
- Jovanovic, Boyan, and Rafael Rob. 1990. "Long Waves and Short Waves: Growth Through Intensive and Extensive Search," *Econometrica*, Vol. 58, No. 6. (Nov., 1990), pp. 1391-1409.
- Kronick, David A. 1976. *A History of Scientific & Technical Periodicals: The Origins and Development of the Scientific and Technical Press 1665-1790*, second edition. Metuchen, N.J.: The Scarecrow Press.
- 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.
- Lipartito, Kenneth. 2003. "Picturephone and the Information Age: The Social Meaning of Failure." *Technology and Culture* 44:1 (January).
- Marsden, Peter V. "Restricted Access in Networks and Models of Power," *American Journal of Sociology* 88:4 (Jan, 1983): 686-717.
- Madey, Greg, Vincent Freeh, and Renee Tynan. 2002. "Understanding OSS as a Self-Organizing Process." <http://opensource.ucc.ie/icse2002/MadeyFreehTynan.pdf>

- Madey, Greg, Vincent Freeh, Renee Tynan, and Chris Hoffman. 2003. "An Analysis of Open Source Software Development Using Social Network." <http://hcs.ucla.edu/lake-arrowhead-2003/madey.html>
- McHugh, Jeanne. 1980. *Alexander Holley and the Makers of Steel*. Johns Hopkins University Press.
- Meyer, Peter B. 2002. "Technological uncertainty and earnings dispersion in U.S. iron work, 1867-1881" <http://techterms.net/ironwork/ironwork.pdf>.
- Mokyr, Joel. 1990. *Lever of Riches: Technological Creativity and Economic Progress*. Oxford University Press.
- Moskowitz, Tomas, and Annette Vissing-Jørgensen. 2001. "The Private Equity Premium Puzzle." Forthcoming, *American Economic Review*.
- Nelson, Richard R. 1982. "The Role of Knowledge in R&D Efficiency." *The Quarterly Journal of Economics*, 97:3 (Aug.), pp 453-470.
- 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.
- Saxenian, Annalee. 1996. *Regional Advantage*. Harvard University Press.
- Scacchi, Walt. 2003. "Understanding Open Source Software Evolution: Applying, Breaking, and Rethinking the Laws of Software Evolution." Institute for Software Research working paper, University of California, Irvine. <http://www.ics.uci.edu/~wscacchi/Papers/New/Understanding-OSS-Evolution.pdf>
- Schrader, Stephan. 1991. "Informal technology transfer between firms: Cooperation through information trading," *Research Policy* 20: 153-170.
- Temin, Peter. 1964. *Iron and Steel in Nineteenth Century America*. MIT Press.
- Tushman, Michael L., and Philip Anderson. 1986. "Technological Discontinuities and Organizational Environments." *Administrative Science Quarterly* 31 (Sept.): 439-465.
- von Hippel, Eric. 1987. "Cooperation between rivals: Informal know-how trading." *Research Policy* 16: 291-302.
- Weber, Steven. 2000. "The Political Economy of Open Source Software" BRIE Working Paper 150. June.