13

Generality, Recombination, and Reuse

Timothy F. Bresnahan

13.1 Motivation and Key Findings

Economists have long noted the benefits to society of recombinant technical change and of general purpose technologies.[1] Recombinant technical change is the reuse of existing innovations in new areas; Schumpeter was probably the first to point out that most technical progress is recombinant. General purpose technologies (GPT) are (a) widely used, (b) capable of ongoing technical improvement, and (c) enable complementary innovation in application sectors (AS).[2] Both recombinant technical change and GPTs involve reuse. From an ex post normative standpoint, reuse creates dynamic social increasing returns to scale and scope.[3] This chapter takes an ex ante positive standpoint and examines the economic incentives and information conditions that lead to original invention of reusable inventions. I emphasize the knowledge available to the inventor, at the time of initial invention, whose work will later be recombined or lead to the emergence of a new

---

Timothy F. Bresnahan is the Landau Professor in Technology and the Economy and professor, by courtesy, of economics in the Graduate School of Business at Stanford University and a research associate of the National Bureau of Economic Research.

I thank Ben Jones, Shane Greenstein, Joel Mokyr, Nathan Rosenberg, Manuel Trajtenberg, Scott Stern, and participants at the NBER Rate and Direction of Technical Change Fiftieth Anniversary preconference and conference for valuable comments.

[1] See, for example, Schumpeter (1939), Nelson and Winter (1982), Weitzman (1998), Romer (1987), Bresnahan and Trajtenberg (1995), and Bresnahan (2010).

[2] See Bresnahan (2010) for the more detailed definitions used in the literature.

[3] I note that the language “increasing returns to scale and scope” implies a normative framework, not a positive one, and similarly that the language “social increasing returns to scale” implies a normative (cooperative) framework rather than a positive (information, incentives, and in this chapter, knowledge) framework. I note also that these benefits assessment frameworks are ex post, that is, recombination, reuse, and generality of purpose are all excellent sources of social gains if they can be achieved.

general purpose technology. Important issues, not well treated in the literature, arise when first inventors do not know of future uses because those uses depend on future invention or on the future creation of new markets and industries.

Recent investigations have deepened our understanding of the logical relationship between reuse and growth theory, and have shown the importance of GPTs in the industrial revolution, the second industrial revolution (in particularly impressive depth), and the information age.[4] Recombination and GPTs can make reuse into a powerful force for economic growth based in increasing returns. Note that this is a normative ex post perspective. Once technologies that can be widely recombined have been invented, once a GPT has been invented and is leading to the further invention of valuable applications, the economy is gaining the benefits of social increasing returns to scale.

In this chapter I focus attention on a new set of corresponding ex ante positive questions about the origins of GPTs and the origins of technologies that will later be recombined. The original invention of a technology that will be widely reused is an important economic event because of the spillovers that flow through reuse.

How, ex ante, are inventors to identify technologies that will be reused or will be general in purpose? Knowledge of what is technically feasible is not sufficient to answer these questions, for an answer depends on future complementary inventions. To make this point sharply, I distinguish between two kinds of knowledge, separating *entrepreneurial knowledge* from the more usual technical and market knowledge. Technical knowledge is a firm's knowledge of its own production possibilities. Market knowledge is what can be observed in existing markets. Entrepreneurial knowledge is, in contrast, knowledge of other firms or industries held in a particular firm or industry. The classical example of entrepreneurial knowledge comes from Hayek (1945). An inventor might know (technically) how to create a new product and yet now know (entrepreneurially) how that product will be used, by whom, and how much value that demand will create. In a decentralized economy, those are all pieces of knowledge (originally) held by others and only learned by the potential inventor at some cost. In the simplest example, a clear engineering plan to build a new mousetrap would be technical knowledge, while knowing ex ante whether the world will beat a path to your door is entrepreneurial knowledge. I extend this concept of entrepreneurial knowledge. The centerpiece of my treatment is that an inventor working in one industry may not know of potential complementary inventions in another industry ex ante.

The point of emphasizing entrepreneurial knowledge is that a market economy typically has highly distributed knowledge. If each agent knows

---

[4] See sources in Bresnahan (2010) and also in Jovanovic and Rousseau (2005).

her own business' invention opportunities and technical needs but not those of other firms or industries—the information requirements needed for a neoclassical economy with price-taking supply—that is distributed knowledge. In this sense, the more distributed is knowledge, the scarcer is entrepreneurial knowledge. This matters for reuse when the knowledge needed to anticipate later uses is not available to an early inventor.

To analyze recombination and GPTs is to consider a world in which there are multiple potential inventors. This leads me to focus on cases in which the economy is decentralized and the resulting potential scarcity of entrepreneurial knowledge is that one potential inventor need not know another potential inventor's circumstances. The inventor of a potential general purpose technology might not, for example, know of the prospects for complementary innovation in applications sectors. Symmetrically, a potential application sector may not know of technical opportunities in what would be, if only it were to be invented, a GPT industry. This kind of scarcity of entrepreneurial knowledge can reduce the ex ante return to innovation.[5]

The second building block of my analysis concerns the way the knowledge state of the economy changes when invention occurs. Suppose once again that ex ante two potential inventors—a GPT inventor and an applications inventor, or an original inventor and a recombiner—do not know of one another's technical possibilities. If, however, one of them has invented something and commercialized it, the other can learn of it. This lessens the scarcity of entrepreneurial knowledge as the second inventor now can look at the first invention and consider whether to make a complementary invention. Of course, the search and information processing need not be costless at this stage. I assume that invention and market presence creates market knowledge, not necessarily complete and perfect market knowledge.

One mechanism by which this might work is if a potential GPT is invented and marketed "on spec," potential applications sector inventors learn of its existence. Entrepreneurial knowledge is then less scarce, and complementary innovation in the AS can be based on market knowledge of the GPT prod-

---

[5] It is a common feature of many economic models of inventions that different inventors have different knowledge. This feature is shared by Schumpeterian models (earlier and later inventors have different knowledge, the later may creatively destroy the earlier), GPT models (GPT and AS have different knowledge needed to work together), recombinant models (ideas become more valuable when combined with other ideas), and standard models of optimal patent policy (early invention and improvement based in different knowledge). The same structure is used in models of organization; each of two agents making complementary innovations has distinct abilities and knowledge.

Another common feature of economic models of invention is the accumulation of a stock of knowledge. Early inventions pave the way for later inventions. Models of quality ladders, for example, assume that each level of quality cannot be invented until after the last level. Models of recombination assume that ideas, once made, can be combined with other ideas in potentially useful ways.

Many of these literatures have been pushed much farther than I attempt here. My goal, however, is to examine the specific problem of scarce entrepreneurial knowledge.

uct. I will call that particular mechanism a "planned initiative." Note that a planned initiative does not require much entrepreneurial knowledge after invention of the GPT. It does require, however, entrepreneurial knowledge ex ante, as the GPT innovator must know what kind of GPT product would appeal to applications sectors. I use "must know" there in an economic sense: the GPT inventor must have a good enough idea of whether AS will follow profitably to invest in a specific technical direction. I will argue that, as a historical matter, planned initiatives are scarce in white-collar work automation (WCA) precisely because this kind of broad-based entrepreneurial knowledge is typically scarce.

When the original problem was difficulties in seeing precise overlaps between technological opportunity and demand needs, early invention and commercialization can create market knowledge of a number of forms. One is that technologists' knowledge of demanders' needs can be converted from scarce entrepreneurial knowledge into widespread market knowledge. Technologists can now learn, by observing what demanders buy, knowledge of what demanders want. A body of demand, once created in a market, can be studied and thus served. An early specific technical solution, even if far from optimal (given all knowledge by both technologist and demanders) can create sufficient market knowledge to enable movement in the direction of optimality. Seeing that a demander is using technology with features G, a technologist can inquire about the marketability of features $G + \Delta g$. If such an inquiry is difficult ex ante, but feasible at the interim stage, valuable market knowledge has been created. Symmetrically, the commercialization of a specific technical product can create knowledge on the part of demanders about what is technically feasible. Demanders could then undertake experiments to see what coinvention works effectively. The results of those experiments are valuable market or technical knowledge; if the results suggest new directions to technologists, they represent an update in the market knowledge of the economy. The fact that demanders needed to undertake experiments can make it very difficult to have complete ex ante entrepreneurial knowledge. A related situation arises when demanders can only understand what a new technology can do by seeing it demonstrated. Their invention of useful applications (which was contingent on the creation of a working prototype technology) can suggest new directions by showing where the overlaps between the technically feasible and the socially desirable.

In a number of historical examples drawn from the computer industry, I examine the case, which I will argue is very important empirically for technical progress in WCA, in which entrepreneurial knowledge is scarce ex ante.[6] We shall see that in an economy with distributed knowledge, overlaps

---

[6] In this regard I follow a long tradition in the analysis of technical change. Like Rosenberg (1996) I emphasize uncertainly and depart from the "linear" model in which science causes technology, which in turn causes application and growth. Yet I also depart from models like that of Acemoglu (2002) in which demand needs are known and directly influence inventors' choices.

between the technically feasible and the socially desirable sets of inventions can be "unknown" in the sense that no individual knows them well enough profitably to direct specific technical investments, and "unknowable" in the sense that either (a) the relevant holders of distributed information need not know one another's technical needs and capabilities with adequate specificity, or (b) detailed good faith discussions among the relevant knowledge holders need not lead to successful communications because the possibility of dual invention is too hypothetical. Initial inventive steps can make the locus of the overlap more known (and more knowable) by converting entrepreneurial knowledge into market knowledge. Since the same industry has launched a number of GPT clusters, it also permits me to examine a number of cases in which entrepreneurial knowledge was less scarce ex ante. The contrast is illuminating about the sources of some of the most important technical advances of the last half century.

13.2 Recombination Model

Economists have already recognized that recombination involves the possibility of knowledge scarcity. Weitzman (1998), in a classic model of recombinant growth, has a model in which the number of "seed" ideas is increasing over time as a result of R&D, and seed ideas can be recombined into potentially valuable inventions. Weitzman's elegant analysis shows first that the combinatorics of mixing and matching an increasing number of ideas can lead to faster- than-exponential expansion of the stock of possible useful inventions (thus easily overcoming diminishing returns). As the number of seed ideas grows, however, the information-processing costs of finding recombinant matches also grow without bound, providing a limit on the growth process. Weitzman's model has no treatment of entrepreneurial knowledge, however. A number of management scholars have taken up the question of search to create recombinant knowledge: a classic study is Fleming (2001), who notes that common knowledge of what technologies are economically related can change over time, and uses the framework of "local" knowledge as related to commercial exploitation of ideas, while "distant" search is exploratory and potentially creates hitherto unforeseen combinations.

An important related notion is that certain kinds of knowledge can come to be science, and that this has important implications for the scope of entrepreneurial knowledge in the economy. Mokyr (2002), for example, makes the important observation that the representation of technical knowledge as science during the industrial revolution in England together with the institutions of open science, lowered the costs of widespread "access" to knowledge. If the solution to the problem of scarce entrepreneurial knowledge is better representation of knowledge, then there is, as Jones (2009) points out, a "burden of knowledge." This suggests an arc of possibility (not unlike the simpler Weitzman arc) in which improving access first improves the ability of the economy to recombine different kinds of knowledge and then creates congestion.

In this section, I model the distinction between different kinds of knowledge related to an invention that may later be recombined, and how the knowledge state of the economy changes ex post its invention. Potential inventors, the only actors in the recombination model, are endowed with technical capabilities and market knowledge, which permit them to make productive inventions at a cost. Potential inventors are also endowed with knowledge about the possible productive applications of their technology. Their entrepreneurial knowledge (or its lack) arises with regard to knowledge about one another.

A simple model can illuminate the economics of entrepreneurial knowledge and recombination. The model is simple in that each potential invention can be recombined either with no other invention or with just one other invention. Potential inventors need not have perfect entrepreneurial knowledge, which in this context means that they do not necessarily know whether their invention can be recombined or, if so, with what.

Begin with a representative invention, called A. The R&D expenditure needed to invent A is r and the return to inventing A if there is no other invention complementary to it is $V(A)$. Any risk, uncertainty, and so forth related to the value of A alone is reflected in $V(A)$.[7] There are a large number of potential inventors of A so that invention will occur if the expected net return to the invention is positive. If there is no possibility of recombination or reuse for A, then the incentive to invent A is given by

$$\pi_A = V(A) - r. \tag{1}$$

Now suppose that there is another invention, B, which can be recombined with A. If both A and B are invented, they can be recombined to create, in addition to the stand-alone values $V(A)$ and $V(B)$, a further recombination value $V(A,B)$.[8] The complementarity behind this additional value is the reason technical change can be recombinant.

If first A and then B have been invented, ex post bargaining or other market transactions between their inventors give the inventor of A a share $\lambda_1 V(A, B)$ of the jointly created value. I am agnostic about how λ is determined, except that I rule out ex ante bargaining because the two complementary inventors may not have heard of one another. The inventor of A

---

[7] I note that $V(A)$ is the return to the inventor. The mechanism by which this return is generated is in the background. If, for example, the inventor of A gets a temporary patent of monopoly on selling A, the total social surplus associated with A will exceed $V(A)$.

[8] The recombinant value could arise because A and B are inventions by a supplier and a customer, or are complements in production or in demand, or because each is a multipart invention and they share a common component. My treatment abstracts away from all those different situations in order to isolate the key problem that arises when inventors of complements do not know of one another.

might get a larger share because a patent regime offers a larger claim to earlier inventors or because the first inventor gets to choose certain market institutions (such as openness) that affect information flows or market power later.[9] The inventor of B will get $\lambda_2 V(A,B)$.[10] Thus, if a potential inventor of A knows that B has been or is about to be invented, the incentive to invent A is given by

$$\pi_A = V(A) - r + C(A,B)\lambda_1. \tag{2}$$

The potential inventors of A, may not, however, know of the pending invention of B or know enough about the characteristics of invention B to assess the prospective increased return from joint invention. The degree to which they *do* know of such things is their entrepreneurial knowledge. I measure entrepreneurial knowledge as a probability assessment, called k, that B can be found by search and is an effective complement for A. Thus the incentive to invent A is

$$\pi_A = V(A) - r + V(A,B)\lambda_1 k. \tag{3}$$

I assume that the invention and marketing of B before the invention of A will improve knowledge about B on the part of potential inventors of A. That is, I assume that after B has been invented and marketed it becomes easier for a potential inventor of A to learn the technical details of B, to make an assessment of the degree of complementarity between B and A, or to design A so that it works well with the B that was actually invented (which may have a higher success rate than designing A to work with a plan of B). This is still entrepreneurial knowledge, but the marketing of B adds some market knowledge to the ex ante guessing and speculation. This higher quality knowledge is represented here by a higher probability assessment that development of A will lead to recombinant value $K > k$.

If there is no complementary technology for A, potential inventors of A may nonetheless think one exists, and have, as a result of this excess optimism, a higher incentive to invent. There is, however, no failure of rational expectations if $k < 1$ for all technologies that are recombinable and no agents with excess optimism. One interpretation of k is the probability that a search for a partner will succeed and an assessment of potential partners'

---

[9] How λ is determined is also pushed to the background. It could arise, as in the models reviewed by Scotchmer (2004), as a result of ex post invention bargaining between the inventors of A and B, each of which has a patent. An alternative mechanism to determine λ is that B sells an input to A and the price of that input, in market context, determines the rent split. I treat these, and other mechanisms to determine λ, as equivalent. It is also not essential that only the synergistic part $V(A,B)$ is subject to bargaining or market division. The claims behind the bargaining reflect not only the formal patent system, but also the openness of the innovation system more generally, the value of first-mover advantages, and so on.

[10] I make no assumption that B gets $(1 - \lambda_1)V(A,B)$. If bargaining or market power is inefficient, as one would expect generally, then the more natural assumption is that B gets less than that.

joint value will lead to a match. This can be less than one for all agents in search of equilibrium.[11]

In summary, k is the measure of the entrepreneurial information available to the inventor whose invention might later be recombined, while K is the measure of the entrepreneurial information of the inventor who might recombine later. After B has been invented, the incentive to invent A rises to

$$\pi_A = V(A) - r + (A,B)\lambda K. \quad (4)$$

A considerable literature has focused on the forces leading, in the language of this chapter, to $K < 1$. An elegant model by Weitzman (1998) illuminates the problem that arises when there are more and more ideas that might be recombined, so that costs of searching among them drive down K endogenously as the overall economy grows more complex and decentralized. The body of work that focuses on "recombinant search" (i.e., search by potentially recombining inventors), focuses on the difficulties in such a search because searchers must cross intellectual or industry boundaries to find and understand potential complements (Fleming 2001). The point of this chapter is that such a search can be even more difficult when the searcher is crossing intellectual or industry boundaries to find and understand potential complements before they have been invented. To search all existing technologies to see which ones offer good opportunities for complementary recombination is one thing; to extend that search to all the as-yet-uninvented technologies that might be a complement and to carefully evaluate their as-yet-undertermined features quite another. Hence my focus on the case where $k < K$.

The novel element here is a distinction between two kinds of knowledge. I distinguish between the technical knowledge of each sector and the entrepreneurial knowledge that has the possibility of creating new markets. Two points about technical knowledge are appropriate here. First, when I simply write $V(A) - r$, I am implicitly assuming good technical knowledge. Second, I am labeling knowledge about the local demand for the invention; that is, what A knows about the probability, demand, and appropriability assessments that lead to value $V(A)$ are all called "technical knowledge." The main point of this is to distinguish it from entrepreneurial knowledge; that is, knowledge about the possible future gains from trade, outside current markets, and connections. I follow Hayek (1945) in making this division between local market or technical knowledge, knowledge about one's own existing business, and entrepreneurial knowledge, knowledge of potential new connections.

The key point about entrepreneurial knowledge is that it only matters before the creation of a new connection. In my framework, once something has been invented and commercialized, knowledge of it is market knowl-

---

[11] I am grateful to Joel Mokyr for useful discussion on this point.

Table 13.1

Agent	Local, technical K	Market K	Entrepreneurial K
Potential *A* inventor	I might invent *A*	You have invented *B*	You might invent *B*
Potential *B* inventor	I might invent *B*	You have invented *A*	You might invent *A*

edge. By that I mean that it depends on what others in the economy are doing, not what they might be doing in a hypothetical future. As a formal matter, this means that invention changes the knowledge state of the economy.

In drawing the distinction between K and k I am implicitly adding a third category of knowledge—market knowledge. If $K > k$ because B has been invented, I call the increase in knowledge about B on the part of potential inventors of A market knowledge. Market knowledge may or may not be perfect, but in table 13.1 I will typically assume that market knowledge about the same outcome is better than entrepreneurial knowledge.

Pulling this together, we have the payoffs relevant to the question of whether recombination will occur. If B has already been invented and marketed, we can focus on the incentives to invent A given market knowledge K. I label this π_{A2} because A is positioned as the second inventor:[12]

$$\pi_{A2} = V(A) - r + V(A,B)\lambda_2 K. \tag{5}$$

There is a symmetric expression for π_{B2}. An idea that is valuable in two uses might be invented first for either of them; it can then be recombined into the other. If B has not yet been invented, however, potential inventors of A will need to rely on their entrepreneurial knowledge to see any benefits of joint invention:

$$\pi_{A1} = V(A) - r + V(A,B)\lambda_1 k, \tag{6}$$

and once again there is a symmetric expression for π_{B1}.

Finally, the order of invention is set exogenously, perhaps by the date at which each stand-alone technology becomes marketable. Without loss of generality (w.l.o.g.), A goes first. One of the many potential inventors of A invents if $\pi_{A1} > 0$. Then, if A has not been invented, one of the potential inventors of B invents if $\pi_{B1} > 0$.[13] If, however, A has been invented, recombinatory technical progress occurs if $\pi_{B2} > 0$. Finally, the opportunity

---

[12] Note that I do not assume that there is some kind of technological hierarchy in which A must be invented before B or vice versa. This assumption is common in the appropriability literature but is not suitable for my purposes. See Scotchmer (2004) for a review of a number of models with this assumption. Technological hierarchy may provide a reason to prefer stronger appropriability for earlier inventors or to oppose openness, an effect omitted from my analysis.

[13] If there were only a single potential inventor of A, that inventor might find it worthwhile to wait for B; with many potential inventors, the possibility of waiting for B is irrelevant in the case $\pi_{A1} > 0$. I am examining a model with such strategic behavior by individual inventors in joint work with Iiro Makinen.

to invent A does not go away, so if B is invented and A was not invented before, that triggers a recombination if $\pi_{B2} > 0$. These conditions determine a (unique) equilibrium as a function of the economic fundamentals.

13.2.1 Social and Private Returns to Invention

For examination of the gap between the social and private rates of invention in this model, I assume

$$V(A) - r < 0; \text{ and } V(B) - r < 0; \qquad (7)$$

$$\text{but } V(A) - r + V(B) - r + V(A,B) > 0,$$

the only interesting case; that is, each stand-alone invention is unprofitable but recombination is profitable.

Consider first the familiar case with no shortage of entrepreneurial knowledge, $K = k = 1$. There is no distinction in this case between π_{A1} and π_{A2} because market and entrepreneurial knowledge are both perfect, and thus both the same; the model is also symmetric. In this case, we can interpret V as a risk-adjusted expected value and interpret (λ_1,λ_2) as the outcome of an ex post bargain between two inventors, limited by their appropriability claims and by imitation. Now, letting A be invented first, the condition for both A and B to be invented is

$$\pi_{A1} = V(A) - r + V(A,B)\lambda_1 > 0 \qquad (8)$$

$$\lambda_{B2} = V(B) - r + V(A,B)\lambda_2 > 0. \qquad (9)$$

Under the assumption of perfect entrepreneurial information, only incentives (λ) matter. If market institutions or patent claims are set up so that one of the λ is too small, then the social rate of return to innovation is less than the private rate of return to innovation. If we force A to invent first (perhaps because the market yielding $V(A)$ opens a century before that yielding $V(B)$) the social return to invention will be less than the private return to invention for A if λ_1 is too small (i.e., [8] fails) and for B if λ_2 is too small ([9] fails). If bargaining is not possible, then the gap between the social and private return to innovation will prevent invention.

Under (7), nondestructive ex ante bargaining, if possible, will always lead a pair of λ, which leads to efficient invention and recombination. Since the two potential inventors know of one another ($K = k = 1$) one can easily suppose that they get together and, for example, form a single firm to internalize the externality of their two inventions; one invents first, and the other recombines into a high-value use. That does not much resemble the "recombination" discussed in the literature, which is part of my point. We now turn to a model in which the opportunity to recombine is unknown ex ante.

13.2.2 Scarce Entrepreneurial Information

Let us now consider a case with the same payoffs and the same timing; that is, joint invention is profitable and the market for A opens first. However, we consider the case with an absence of entrepreneurial information ($k = 0$) together with excellent market information ($K = 1$). Under these assumptions, the condition for both to invent is

$$\pi_{A1} = V(A) - r + V(A,B)\lambda_1 k > 0 \Leftrightarrow V(A) - r > 0$$

$$\lambda_{B2} = V(B) - r + V(A,B)\lambda_2 K > 0 \Leftrightarrow V(B) - r + V(A, B)\lambda_2 > 0.$$

The second condition, recombination by an inventor of B, will be satisfied for some admissible λ_2. The first condition, however, cannot be satisfied when only joint invention is economic (7). Reversing the order or having the potential inventors have the opportunity to invent simultaneously does not help. It is easy to see there will be no first invention under (7). The problem here is that valuable invention is not undertaken because it only becomes sufficiently valuable in the information state—unknown to an original inventor—that it will be later recombined. The fact that invention will create that information state ($K = 1$) is not helpful when the information does not exist.

Increasing original inventors' share of eventual returns by raising λ_1 does not change their incentives to invent, because λ_1 is multiplied by zero. Since the original inventor does not know about the future recombination that may create recombinant value ($k = 0$), giving them a larger share of the returns from future recombination is pushing on a rope. Changing from open innovation systems to closed, or allocating stronger patent claims to earlier innovators as a strategy to increase λ_1 is ineffective, and, to the extent it decreases λ_2, dysfunctional. The later, recombining inventor acts at a time of better information, so the decrease in their incentive to invent is far worse than the benefit to A.

This example, while extreme, reveals the importance of entrepreneurial knowledge. An invention that will gain value from later being recombined will, more generally, not have adequate invention incentives if the first inventor does not know about the potential recombination. Note that this effect does not depend on there being anything odd about the first inventor's knowledge of her own business or her own market. She can be perfectly rational, perfectly foresighted, understand all technical possibilities without regard to whether they involve a conceptual breakthrough or not, and so forth. The key assumption is one of limited entrepreneurial knowledge in the sense that knowledge is held in a distributed way (i.e., that she does not know about future technical possibilities in another business where her invention might be recombined).

In this case, the private return to innovation is below the social return to innovation if we evaluate returns using the ex post knowledge, or to put it another way, using the standard first-best assumption that we the analysts have all of the information in the economy.

This kind of scarce entrepreneurial knowledge raises the social return to innovation above the private return. Indeed, whenever we see recombination, it is reasonable to suspect that earlier entrepreneurial information about the then-future recombination was scarce. The private incentive of the original inventor to invent fell below what we now know, using ex post knowledge, was the social incentive. But this argument must be treated very carefully. The high "social return to innovation" of the first innovation can be calculated only by using all the information in the economy, not the information available to any inventor. Nor can conventional incentives (claims, market positions, etc.) raise the private return up to the social return.

Bargaining among the two inventors is not a solution. Search by potential inventors of A has either not led them to locate potential inventors of B, or has not convinced them adequately of the proposition that B might be a complement to act on it.

13.2.3 Comparative Statics

Each of the two first examples was extreme. More generally, even when we let both k, K, and λ be arbitrary, we get the result that, the more important is low k as a source of poor returns to innovation that might be recombined later, the weaker are increases in λ_1 as a mechanism to overcome it. Similarly, the larger is K relative to k, the greater is the improvement in knowledge about potential recombination, and thus the greater the advantage of giving incentives to later inventors (λ_2). Neither of these points turns on the extremity of the examples. Another comparative statics point that would arise in a more fully articulated model is that rather than not being invented at all, a first invention of a recombinant pair might be invented with too low a probability (if, e.g., r is a random variable) or at too late a date (if, e.g., V are rising because the economy is growing or r is falling because of technical progress elsewhere). In my historical examples, I will make obvious extensions like these without a formal model.

13.2.4 Remarks

The novel idea in this section is that the invention and commercialization of a technology depends on entrepreneurial knowledge and creates market knowledge. This puts recombination in a new light. In a decentralized economy, the ex ante perception that a particular invention might later be recombined is entrepreneurial knowledge. Scarcity of entrepreneurial knowledge ex ante, like the more familiar problems of weak appropriability or scarce technical knowledge, limits incentives to innovate. Evaluating either the private or the social rate of return to invention using all of the decentralized knowledge that exists in the economy would reveal the positive returns flowing from recombination. The problem in the case of scarce entrepreneurial knowledge is that no one knows enough to make the calculation.[14]

---

[14] Hayek (1945, 519–21): "The peculiar character of the problem of a rational economic order is determined precisely by the fact that the knowledge of the circumstances of which

To the extent scarce entrepreneurial knowledge is a source of deviations of private from social returns to innovation, it suggests narrow patents or open systems rather than giving original inventors broader claims. Giving broad claims can be actively counterproductive (above and beyond not being productive ex ante) if the rights given to original inventors are broad enough to encompass unforeseen recombination. They limit the incentives of later inventors, who work in a better knowledge environment.

If the problem in innovation is scarce entrepreneurial knowledge, one could think that the solution is teaching everyone what everyone else knows. If that means lowering the costs of storing, retrieving, and communicating knowledge, reducing the possibility that distributed knowledge is a bottleneck, it makes excellent sense. For example, the available stock of knowledge in the economy might be partially codified into science, and access costs to that science could be lowered. This creates a widespread knowledge asset, reducing the degree to which technical knowledge is local. Of course, as the total volume of knowledge rises, the costs of information processing can make this less effective.

It is worth pointing out that all of these normative ideas, however valuable within their scope, may be of limited relevance to the economic problem of an initial invention that later is reused. Making knowledge that already exists easy to retrieve broadly is a good thing; making knowledge that does not yet exist or which is not yet known to be useful to anyone easy to retrieve risks clogging the system. Further, there are excellent reasons, related to the day-to-day functioning of the economy, why much commercial knowledge is decentralized, so it may simply not be cost-effective to have everyone know everyone else's business well enough to know exactly what everyone else might create. In short, the shortage of entrepreneurial knowledge in the economy may be a social cost.

Indeed, I shall argue in my historical section later that we should understand the entrepreneurial-knowledge shortfalls that bottlenecked some very important late twentieth-century GPTs were, in fact, social costs. My argument there is grounded in specific historical detail, of course, but the general analytical point is clear.

13.3 The Founding of GPT Clusters

I now turn to the founding of GPT clusters. A GPT cluster consists of a GPT and several applications sectors. The underlying model of a GPT

---

we must make use never exists in concentrated or integrated form, but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess. The economic problem of society is thus not merely a problem of how to allocate 'given' resources—if 'given' is taken to mean given to a single mind which deliberately solves the problem set by these 'data.' It is rather a problem of how to secure the best use of resources known to any of the members of society, for ends whose relative importance only these individuals know. Or, to put it briefly, it is a problem of the utilization of knowledge not given to anyone in its totality."

cluster shares one crucial feature with the model of recombination in the last section: there is complementarity between invention of a GPT and invention in each applications sector. The most important difference is that a GPT has more than one potential AS partner.

Thus the simplest GPT cluster consists of three potential inventions, A_1, A_2, and G. Each of them costs r_a to invent, and each creates a stand-alone value $V(a)$, $a \in \{A_1, A_2\}$, and $V(G)$. There is also an innovative complementarity between each of the applications and G, so that a further value is created if either both A_1 and G are invented or if both A_2 and G are invented. Call this value $V(a,G)$. No (direct) innovative complementarity exists between A_1 and A_2, though as has been noted in many contexts potential inventors of these two technologies have a common interest in G.

By assembling all the distributed knowledge, we know (correctly) that one of these technologies is a GPT (G) and that the other two are potential applications sectors for it. Potential inventors, however, need not know this ex ante. The notation for who knows what is now necessarily more complex: I denote entrepreneurial knowledge once again by k; now k_a^G refers to knowledge held about G by potential inventors $a \in \{A_1, A_2\}$ while k_G^a refers to knowledge held about a by potential inventors of G. After any technology has been invented and marketed, market knowledge is created. Once again I use K to denote this, and the notation K_k^j denotes knowledge held by potential inventors of technology k about technology j after j has been invented and marketed As in the previous section, the obvious assumption is $0 \le k_k^j \le K_k^j \le 1$ for all pairs j, k. Once again I will denote the share of the complementary return that go to each of the two parties (G, an a) by λ.

The market relationships between a potential GPT and potential applications sectors before and after innovation will influence k and K. In one case, G is a process component that can be used in production in a. Then we should expect k to be low and particularly so if potential inventors of G are already, preinvention, suppliers of a. If G instead is an enabling technology, such as a tool to permit inventions in a, we should expect k to be higher, and particularly so if the "coinventions" in a are itself hard to foresee. A G that primarily enables radical coinventions will have lower k than one that enables nondisruptive ones, and so on. Some cases of GPT platforms are likely to have lower k, or to call for a wider span of k. If applications share customers, and if customers must select G (one kind of platform market), a potential inventor A_1 may need entrepreneurial knowledge not only about G but about the customers A_2 may attract to G.

13.3.1 No Invention

Scarce entrepreneurial information or weak incentives can lead the private rate of return to be less than the social rate of return (the latter assessed using all the information in the economy). In particular, either low k or low λ_1 can lead to failures of the condition to invent:

(10) $$0 > V(G) + \sum_{a} V(a,G)\lambda_1 k_G^a - r_g$$

(11) $$0 > V(a) + V(a,G)\lambda_1 k_a^G - r_a \forall a.$$

13.3.2 Planned Initiatives

There is a natural tendency to think of GPTs in a hierarchical way. Someone invents a GPT, offers it to potential users, and induces applications sector investment in complements. The GPT inventor might also design a "local" patent or copyright regime that applies to *A* that work with *G*. In this section, I call such a path to the invention of an entire GPT cluster a "planned initiative" and point out that a successful planned initiative turns on the entrepreneurial knowledge of the firm designing the practical GPT product.

A planned initiative is the only equilibrium if a potential inventor of *G* has an incentive to invent and applications sectors have an incentive to follow but not to lead:

(12) $$V(G) + \sum_{a} V(a,G)\lambda_1 k_G^a - r_g > 0$$

(13) $$V(a) + V(a,G)\lambda_2 K_a^G - r_a > 0 > V(a) + V(a,G)\lambda_1 k_a^G - r_a \forall a.$$

This condition states that no potential GPT inventor has an incentive to invent as a planned initiative, anticipating follow-ons by *a*, and it succeeds in getting some complementarity value if any *a* follows, while generality is achieved if more than one *a* follows. The incentive for the GPT to be invented first need not involve contractual understandings with the *A* sectors. Instead, it may be undertaken "on spec" with the k_G^a measuring the probability assessment on the part of potential inventors of *G* that there will be an application of type *a*. For a planned initiative to succeed, the key entrepreneurial knowledge is that of the GPT or platform innovator. The innovator must have a wide enough knowledge of potential applications to assess the likelihood of success. In a planned initiative, the applications sectors come second, and thus need not have entrepreneurial knowledge of *G*, as they can see *G* in the marketplace.

When ex ante bargaining is feasible and entrepreneurial information is good, another form of planned initiative can arise in which a GPT inventor and one or more early inventors of applications set up incentives for later applications inventors.

13.3.3 Technological Convergence

The other extreme form of equilibrium in the GPT case is technological convergence (Rosenberg 1963). This denotes the case in which the *A* are invented first and only later does a general purpose technology arise. Whereas in a planned initiative, the general leads the specific, under technological convergence, specific solutions emerge first and are later generalized. The conditions for technological convergence to be the unique form of equilibrium are

$$(14) \qquad V(G)+\sum_a V(a,G)\lambda_2 K_G^a - r_g > 0 > V(G)+\sum_a V(a,G)\lambda_1 k_G^a - r_g$$

$$(15) \qquad V(a) + V(a,G)\lambda_1 k_a^G - r_a > 0 \;\forall\, a.$$

As Rosenberg (1963) points out, one attractive theory of technological convergence is that no one knows ex ante that there are common elements of the production process in A_1 and A_2. There is no technological reason for the general to be invented before the specific, especially if the specific has the goad of necessity. However, after each industry has improved its production process separately, the common elements can be seen more easily ($K_G^a > k_G^a$ in the notation of this chapter). At that point, their common technological elements can be turned into a common technological component supplied by a GPT industry. Invention of the general takes the form of abstracting from the specific.

The case of technological convergence brings out an element of GPTs that many have noted, which is the (social) increasing returns to scale that can be obtained by sharing a common, general, technical input across many applications sectors. This can be salient to the conditions that prevent emergence of a planned initiative. Consider the case in which $k_G^{A_1} > k_G^{A_2}$ and in which the profitability of a GPT turns on it being used widely; that is, on finding all the specific complementary investments in different applications. Then planned initiative might not arise because condition (12) fails, not because there is no idea that the technology inherent in a GPT is useful, but because full range of complementary investments that are necessary for a general solution to be economic are not yet visible.

Note that it is not possible to change only λ and switch conditions in which a planned initiative is the only possible equilibrium to conditions in which technological convergence is the only possible equilibrium. It is as straightforward as possible to obtain such a switch by changing k.

13.3.4 Inversion

In the simple three-inventor model, let (w.l.o.g.) $V(A_1, G) > V(A_2, G)$. In this model an inversion is the invention of A_2 first, followed by G, then followed by A_1. I call this form of equilibrium an inversion because the order of discovery of applications for the GPT is the opposite of the order suggested by valuation. The conditions for an inversion are

$$(16) \qquad 0 < V(A_2) + V(A_2, G)\lambda_1 k_a^G - r_a$$

$$(17) \qquad 0 > V(G) + \sum_a V(a,G)\lambda_1 k_G^a - r_g$$

$$(18) \qquad 0 > V(A_1) + V(A_1, G)\lambda_1 k_a^G - r_a$$

$$(19)\qquad 0 < V(G) + V(A_2,G)\lambda_2 K_G^{A_2} + V(A_1,G)\lambda_1 k_G^{A_1} - r_g$$

$$(20)\qquad 0 < V(A_1) + V(A_1,G)\lambda_2 K_a^G - r_a.$$

The first two inequalities are the core distinctions between an inversion and a planned initiative or technological convergence. Inequality (16) says that an applications sector invents before any G is invented. This is like the condition for first invention in technological convergence, except that it only holds for a single sector—in the case of an inversion, a low-value sector. Inequality (17) is the opposite of the G-invention condition in a planned initiative; here, no potential inventor of a G can be adequately sure of complementary applications development to invent.

The essential feature of an inversion is thus that incomplete entrepreneurial information block joint invention of G with the most valuable application but not with other applications. This looks odd from an ex post perspective but not from an ex ante one.

To get inversion as a likely market form, we need some force that creates a negative correlation in the cross section of a sectors between $V(a, G)$ and k_a^G. There are, of course, ways to make this true. If high value applications sectors are the ones, for example, which need to experiment to take advantage of a new G capability, that would imply such a negative correlation and thus the inversion. Thinking we need a "negative correlation," however, turns on using an ex post perspective, which uses knowledge no potential inventor has ex ante. One good ex ante comparison of the conditions for inversion is to the conditions for technological convergence. If the different applications sectors are thinking about their own businesses, the key assumption behind an inversion is that only one sector invents. Neither that sector nor the applications sector that does not invent knows the relationship of complementarity between their innovation and a new technology to be invented in the future.

Another way to say this same point is that inversions tend to arise when there is a gap between social and private returns to innovation looking at the GPT and its highest value application. This also makes it clear why inversions can lead to the creation of great value. Inequality (19) holds if the invention of A_2 creates market knowledge $K_G^{A_2}$ for potential inventors in G that leads them to invent (this is much like the condition for a GPT to invent in technological convergence). Inequality (20) means that the invention of G creates market knowledge, which leads to further application.

It is that last step that I call an *acceleration.* There is an acceleration in value creation as additional sectors invent. What is going on in the acceleration is the release of the market from the bottleneck that held the private rate of return to invention below the social rate. To the extent that lack of entrepreneurial information can create a low private rate of return to invention, the acceleration in value creation is unsurprising.

The triggering event for the acceleration is the *decentralization* of invention that follows from the creation of market knowledge. In an inversion, no single agent knows enough to coordinate, and the ex ante costs of search are too high to make economic coordination possible. However, the early inventions create market knowledge, which raises the private return to other inventors. The central point here is that the decentralization of invention is part of an inversion because of the assumption of distributed knowledge.

An inversion is a market work-around to lack of entrepreneurial knowledge about the value of coordination between potential inventors of G and of A_1. The generality of G is an important assumption here. Looking only at G and of A_1's lack of entrepreneurial knowledge blocks valuable coordination of invention. The generality of use of G permits invention despite this. Of course, this is not a first-best argument. The market work-around cannot occur unless the less valuable applications are still valuable enough to pay for inventing G. Nonetheless, the possibility of accelerating value creation in the late stages of an inversion is valuable. And it is important to point out that, in conditions of limited entrepreneurial knowledge, this market work-around is feasible, where contracting to overcome the coordination problem is not feasible because of the distributed knowledge. I call this a *market* work-around to contrast it with much *antimarket* thinking about the origins of platforms and of GPT industries, focused on contracting and bargaining.

A planned initiative is not the only path to invention of a GPT. Innovation in a number of important GPTs has followed a "circuitous route." I define a circuitous route as having three characteristics: (1) inversion, (2) decentralization, and (3) acceleration. In this section, I show a model that makes definition of all three elements precise. (1) Inversion: The first invention leading to creation of a market in the GPT has a narrow and specific purpose serving a moderately valuable use. The economic motivation of the original invention does not include either generality of purpose or more valuable uses than its narrow and specific purpose. (The word "economic" here is important. Many inventors hope and anticipate that their invention will be generally useful, and it is important for causal arguments that this does not always lead to investment in their invention.) (2) Decentralization: A series of innovations, arising from a number of sources, leads to the successful exploitation of the ex post more valuable uses. Key steps in this sequence of innovations are not coordinated ex ante; instead, early innovations create knowledge about markets that informs later innovators. (3) Acceleration: Once it is known that the "GPT" is general, the positive feedback associated with social increasing returns to scale raise the returns to invention of improvements to the GPT and coinvention of applications.

13.3.5 Multiple Variants of G

Another point can be made in the standard model in which the AS are symmetrical in value—but here, various with regard to entrepreneurial information. Suppose that for each A, there are two potential ways to create new value. One is a compromise, specific to the sector and involving invention of A and $g = \gamma(a)$. The other is an efficient general to all sectors and involves invention of A and G. To capture "compromise" and "efficient" assume that $V(A,G) > V(A, \gamma(A)) > V(A, g)$ for all other g, notably including $\gamma(b)$. However, $r_G > r_{\gamma(a)}$ for all a, so the generality is expensive.

Add the assumption that entrepreneurial knowledge about $\gamma(a)$ is good, but that potential inventors of G have good entrepreneurial knowledge about applications in only ρ proportion of cases in the sense that

$$V(G) + \sum_a V(a,G)\lambda_1 k_G^a - r_g = V(G) + \rho \sum_a V(a,G)\lambda_1 - r_g.$$

Note that this condition has the advantages and the disadvantages of scope. The advantage of wide scope of applicability is that the fixed cost r_G is spread over many AS. The corresponding disadvantage arises when entrepreneurial information is scarce, for then potential GPT inventors may not know of the specific needs of their potential customers. In the case where ρ is small then the absence of entrepreneurial information about broad opportunities makes invention of a GPT on spec uneconomic. Of course, if this expression is positive, G is invented in a planned initiative and all is well. But what if it is not?

Let us assume that with an A to invent using $\gamma(A)$ (recall they have perfect entrepreneurial information) the comparable condition for invention is

$$V(a) + V(\gamma(a)) + V(a,\gamma(a)) - r_a - r_{\gamma(a)} > 0. \tag{21}$$

Assume that the proportion of applications sectors for which this will hold is Ψ and the proportion of applications sectors for which this will hold and that are entrepreneurially known to potential inventors of G is $\rho\Psi$. Then a general GPT will be invented after a first round of invention of A and $\gamma(A)$ in some sectors if

$$V(G) + \rho(1 - \Psi) \sum_a V(a,G)\lambda_1 + \Psi \sum_a V(a,G)\lambda_2 - r_g > 0. \tag{22}$$

This can be substantially larger than the condition for original invention of a GPT if there is enough opportunity to create local alternatives. It is worth noting that strong patent rights for these alternatives (enough to lower λ_2) can still prevent emergence of a general purpose technology.

13.3.6 Remarks

In this section I have constructed a model with the simplest structure that explains inversion, one built around limited entrepreneurial knowledge. Inversion is an odd enough phenomenon that it calls for adding something to the model. An added benefit is that the model predicts decentralization and acceleration. It explains why, in the case of a GPT, a market work-around is available to deal with bottlenecks caused by entrepreneurial knowledge scarcity. How important these phenomena are can only be investigated by close historical examination of the knowledge state of the economy at different stages of invention. I will show that these ideas, especially the replacement of scarce entrepreneurial knowledge with excellent market knowledge, mattered for the rate and direction of technical change.

13.4 Historical Examples

I now turn to six historical inventions of important GPTs, all within computer and information technology. Three of them are the three most important (so far!) computer GPT clusters for white-collar automation (WCA). These are (1) business data processing based on computers, (2) personal computing, and (3) the widespread use of the Internet and the World Wide Web (WWW). These three GPT clusters have included—but not begun with—a wide range of WCA applications respectively in (1) enterprise computing, (2) personal productivity computing, and (3) electronic commerce, communication, and content. The third recombined the first two (and a number of other technologies) and its applications have considerably expanded the demand for them. I also study the founding of two other important GPT clusters within the same technology category, with very different conditions of entrepreneurial knowledge. These are (4) the computer itself, (5) the minicomputer, and (6) the smartphone. At the beginning of each of those segments, an innovator had the entrepreneurial knowledge to see the overlaps between the feasible and the valuable. The contrast to my first three examples is instructive.

I study the creation of information technology GPTs for three related reasons. First, these are, particularly in their application to WCA, among the most important contemporary technologies. Second, there is a large body of careful historical studies of invention in this industry.[15] My brief treatment builds on these, and a focus on a novel historical question. Specifically, I focus on the knowledge state of the economy before markets were founded. Earlier studies have been strong on what specific firms or individuals knew and thought, laying a very strong basis for my work. Each of these reasons to study information technology GPTs is standard and simple. Each of these three began in an inversion, following, at least for a while, a circuitous route to its highest value applications.

I also turn to information technology GPTs because, at least in WCA applications, entrepreneurial knowledge has often been scarce. In particular, it has been difficult to see overlaps between the technically feasible and the

---

[15] I draw heavily on Aspray and Campbell-Kelly (2004), Ceruzzi (1998), Freiberger and Swaine (2000), Langlois and Robertson (1992), on Usselman (1993), and on my collaborations with Shane Greenstein (1996) and Franco Malerba (1998). In some of the historical episodes I draw on new primary sources.

valuable in application. This has been noted in the past as a source of failure of cutting edge applications, a source of the slow diffusion of valuable new technologies, and an explanation of firm success based on marketing capabilities.[16] Thus I am, to a considerable degree, looking for the problem of scarce entrepreneurial knowledge where I expect to find it. That creates obvious problems, which I overcome by looking at GPT clusters based on the same broad technology area founded in conditions of better entrepreneurial knowledge.

Another advantage of these historical examples is that they help sharpen both the concept of entrepreneurial knowledge and its economic role. Conceptually, entrepreneurial knowledge must be (a) specific enough to guide investment in new technology and (b) connected enough to create a market. Grace Hopper's distinction between thinking a computer (or other new invention) was a good idea and actually building a computer that solves a problem captures much of this.[17] I would add the economist's point to that; only ideas specific enough to draw investment resources are K. We shall see that the distinction between broad general knowledge that some invention in a wide technological area might be useful and knowing a direction for technical progress that might well serve an identified user need is crucial for drawing investment resources.

13.4.1 Entrepreneurial Knowledge Scarcity and Market Work-Arounds at Industry Founding

Because of a scarcity of entrepreneurial knowledge linking an important technology to its most valuable use, one of the twentieth century's most valuable GPTs, business data processing, was invented in an inversion. The key shortage of entrepreneurial knowledge arose here: It was difficult to see, ex ante, the overlaps between what was technically feasible and the most valuable uses. What was clearly technically feasible was the computer; how to make a computer valuable in business was not obvious, especially not to those who best understood business data processing. The overlap between the technically feasible and the valuable in use became more visible at an interim stage, after the invention of general purpose computers to meet significant, but lesser, demand needs. To understand this more clearly, I examine the invention of the computer itself, the founding of the business data processing industry, and the founding of the minicomputer industry.

---

[16] This has been well-documented in the writings of industry insiders (e.g. Gates 1995, see note 32). Shane Greenstein and I pointed out that the importance of marketing capabilities at the firm level has historically been far greater in the commercial (mainframe, PC, smartphone) segments than in the technical ones (minicomputer, engineering workstation). As we shall see, this is related to the relative importance of incomplete entrepreneurial knowledge in the commercial sectors.

[17] Admiral Hopper was the inventer of the compiler.

13.4.2 The Computer

Much of the foundational engineering advances that constitute invention of a general purpose electronic computer were undertaken in the 1930s and 1940s, though it would take a large number of improvements and extensions over more than a half century to create all of the technologies now supporting white-collar automation. The same half century contained a looming growth bottleneck for the rich countries. Automation of physical processes and of blue-collar work in many industries (e.g., in manufacturing), was very successful but was, over the next half century or so, destined to be subject to diminishing returns. One thing clearly needed for further growth was technical progress in WCA.

Today, we all know that one group of uses of electronic computing was going to be business data processing for automation and product quality improvement in service industries and for the white-collar functions of all industries. Today, ex post it is obvious that computer-supported business data processing is a valuable overlap between technological opportunity and demand need. As ex ante entrepreneurial knowledge, it was far less obvious. To be sure, there was a great deal of excitement about the prospects for computers, largely among scientists and engineers interested in *calculation* (military or civilian).[18]

Much of the specific progress that was made in computers in the late 1930s and in the 1940s was to make machines that could compute; that is, do arithmetic calculations. Specifically, they were invented by scientists and engineers to support scientific and engineering calculations, frequently supported by military funding. Very important examples include the work, funded by the Army, of Eckert and Mauchly at Penn, and the work of physicists and mathematicians recruited to work on atomic weapons projects, notably John von Neumann. The scientific and engineering calculations they wanted to undertake included some that were numerically difficult, such as making artillery tables, and others that were both conceptually deep and numerically difficult, such as the calculations needed to design the H-Bomb, which involved understanding some of the deepest mathematics and physics ever conceived. From the perspective of entrepreneurial knowledge, however, it is entirely correct to assume a large k_a^G—the *relationship* of the desired calculations to a machine that could do calculations was, unlike the calculations themselves, not complex. That relationship is *entrepreneurial knowledge.* This is how the inversion started. One potential group of *A*s, scientific computing, had very good entrepreneurial knowledge.

It is also helpful to locate technical knowledge and entrepreneurial knowledge together. Scientists and engineers were also well set up to understand

---

[18] There were also widespread forecasts that computers would be useful for everything. This is not the same as entrepreneurial knowledge that guided investment.

the technical requirements of an electronic calculating machine itself. They could see, once the problem of creating a machine to undertake calculations was set, paths to making such a machine. Of course, it was extremely helpful that some of the goals of calculations were obviously beneficial in a military sense, so the scientists and engineers were often well-funded. This was an example of particularly good entrepreneurial knowledge about the value of a new tool, electronic computers, held by people with the knowledge to make it—physical scientists and engineers.[19]

Many, many inventors have claimed to be the first in some aspect of the electronic computer, the stored program computer, and so forth. This includes a claim from IBM, later to be the most successful electronic data processing firm using electronic computers, related to their joint work with Howard Aiken of Harvard in the late 1930s and during the war. This claim is important to the present inquiry because IBM was (like others) already engaged in business data processing in the 1930s. However, IBM did not have the requisite entrepreneurial information to invest in electronic computing for business data processing. Like its competitors, IBM was investing overwhelmingly in research and development of mechanical and electromechanical technologies, not in digital computers. The Aiken project at Harvard was to create a machine that could do calculations in physics (Aiken's department). Aiken was looking for a calculation firm, and turned to IBM only after Monroe (the calculator company) turned him down. The Aiken project used IBM's existing electromechanical technologies, so the direction of technical progress here represented the recombination of IBM technologies with scientists' needs, not the other way. The point here is absolutely not to belittle the inventiveness and foresight of this project. Instead, the point is to say what this project was not: an investment by IBM in technologies to be useful in business data processing. It was only much later, as we shall see following, that IBM turned to the overlap between electronic computing and business data processing.[20]

The core distinction here between scientific calculation and business data processing at the earliest stages concerns the presence of actionable entrepreneurial knowledge. Military demand for scientific calculation had it; Aspray and Campbell-Kelly (2004) correctly begin their chapter entitled "inventing the computer" by saying "World War II was a scientific war." In contrast, their next chapter, entitled "The Computer Becomes a Business Machine," begins with a story about Thomas Watson (sr.) of IBM. After about 1951, IBM became very aware of the potential of the computer as

---

[19] It is, of course, not true of scientific and engineering tools in general. Those are often built in interdisciplinary teams where one knows the purpose and the other the methods. Entrepreneurial knowledge is needed for that.

[20] IBM did not take up the burst of technical opportunity that arose in World War II; it was not until the Korean War that "government sponsored competition" prompted IBM to move into computing.

a business machine and played a central role in its reconstruction “to be a data-processing machine rather than a mathematical instrument” (Aspray and Campbell-Kelly 2004, 106). How this reconstruction was undertaken is well understood: how it was enabled is an important part of the inversion that created the data-processing machine.

13.4.3 Openness

The sense in which the scientists and engineers invented a GPT was that they invented and improved *tools* that they could use in scientific and engineering calculations. As is the habit of scientists, they designed the tool to be general calculating engines. A scientist does not make a tool general because she or he foresees all its uses; on the contrary, generality is often motivated by a sense that others may take up the tool for their own uses. The essential role of science here was the openness with which the tool was delivered to the rest of the economy, including other scientific and engineering disciplines, and ultimately to unrelated commercial application.[21]

This tool turned out to be suitable for recombination outside the range of science and engineering. That recombination led to a very large spillover from the scientific sector to the rest of the economy (to which we shall turn in a moment) but the spillover did not flow through application of the science itself. The essential role played by the scientific-ness of the original inventors in the spillover process was not the new scientific knowledge itself. The spillout was the recombination of an input into science. This is not “the commercialization of science” as often understood, but the beneficial effects of scientific openness in widespread dissemination of a tool.[22] The organizational structures and values that supported openness, generality, and disclosure, which exist in scientific communities, to be sure, but also in some other invention communities, can form very important parts of a market work-around when the linear path is blocked by lack of entrepreneurial knowledge. In this case, IBM’s λ_2 would be quite large, and the original inventors of many critical computing technologies did not command much of a λ_1 once the computer was recombined into business data processing.

Once the computer had been invented and was being applied to an every widening circle of computations, the knowledge state of the economy changed. What had been scarce k became very widely held K. Many people could now see the possibility of the general purpose computer as a business

---

[21] It was thought for a time that Eckert and Mauchly had a patent as a result of the Electronic Numerical Integrator and Computer (ENIAC), but this turned out to be incorrect. Some say the commercially-oriented Eckert and Mauchly invented the stored program computer, others say it was John von Neuman. There is no doubt, however, that it was von Neuman who sought to have the concept and engineering of the stored program computer available to all.

[22] The discovery and associated inventions of the semiconductor effect, the transistor, and the integrated circuit were an extremely important spillout from science to the computer industry, and were very much the commercialization of science.

tool, at least in applications that were obviously computational, such as accounting, finance, and some operations management tasks like inventory control. To be sure, the electronic computer would have to overcome serious disadvantages relative to electromechanical devices, such as low reliability. That, however, could be conceptualized as a technical/engineering problem.

Perhaps more importantly, once the technical knowledge about the computation itself was made open, it could be combined with other knowledge about business data processing. This was a far more difficult problem than scientific calculation. Most managerial applications of business data processing have a very complex relationship between the business logic of the application and the technical capabilities of computer hardware and software. The simplest are accounting and finance and even they have a more complex interface with calculation than do typical scientific or engineering calculations. This very complex problem was, however, partly solved by the invention of the electronic computer as a mathematical engine. A decentralized path of invention could take advantage of the widely distributed knowledge in the economy, and now firms with knowledge of business data processing entered the picture in a very important way.

It is a mistake, a very common mistake, to think that the only entrepreneurial information problem at an early stage is a shortage of "vision" on the part of "visionaries"—that is, individuals or firms who foresee the future. This misses a central important point about entrepreneurial knowledge. Market economies can, with the help of enough openness, achieve breakthroughs that were unforeseeable to any individual because knowledge was widely distributed. Of course, those breakthroughs that arise through an inversion come later than they would have if there had been a single individual with all the knowledge of both technical possibility and demand needs. The distributed state of ex ante knowledge is a social cost, but at what a high rate and in what an excellent direction technical progress can proceed ex post an inversion. That improvement arises from opportunity pent up while the social return to invention is above the private return, opportunity unleashed by the changing knowledge state of the economy.

13.5 A Planned Initiative Succeeds

Once an inversion is completed, the newly created information about technical progress may lead, through decentralization, to recombinant invention by distinct inventors than those who participated in the original inversion. Those new inventions can lead to an acceleration, completing the circuitous route to the founding of a market.

A wide number of firms, with an extremely wide range of knowledge bases and capabilities, entered a race to be the leading computer vendor in business data processing. IBM, though its technical knowledge base lay in mechanical and electromechanical business data processing, won this race.[23] IBM took advantage of newly public knowledge about computers, its own existing knowledge about the needs of business data processing, and undertook significant recombination.

Open knowledge about the electronic computer did not just benefit IBM. The openness created a large number of recombinatory experiments in competition with one another. No one knew exactly what a business data processing computer looked like even after they saw a successful scientific computer. The competitive experimentation race to establish a successful business data processing business around the mainframe computer worked well in such a knowledge-challenged environment.

In the case of business data processing there was still a great deal of invention to be undertaken in computers themselves and in their commercial applications to build a complete GPT cluster. What is quite interesting about those next steps is that they took a radically different form: IBM undertook a planned initiative to construct a GPT cluster centered on the mainframe computer and induced customers, primarily large firms, to create applications. That planned initiative won a competitive race among a number of distinct business data processing firms that ended with an IBM standard.[24]

IBM went to work to create the general purpose components that could be used by its corporate customers to build applications. IBM also built a very good computer design and engineering technical capability, though IBM was rarely the technical leader in computers, narrowly understood. Yet IBM offered a complete set of complementary general-purpose inputs, including hardware, software, storage, and other peripherals that reflected its knowledge of the kind of problems its customers were trying to solve. Further, IBM put in place an organizational support system that let its customers lower the risks of undertaking experiments in the applications of computers—this is a general purpose complement unmatched by any significant competitor worldwide. The creation of the IBM mainframe standard was an example of how a planned initiative can build a GPT cluster. To underscore the key point here, once IBM understood the technical prospects for electronic computing reasonably well, that single firm had the entrepreneurial knowledge to undertake a planned initiative. It combined preexisting knowledge of its customers' needs with new, generally available knowledge about what was technically feasible.

Of course, there was continuing feedback between technical knowledge and knowledge of user needs in computing for decades after this. There was a dramatically high rate of technical progress in computing, even if we think

---

[23] See Bresnahan-Malerba on the nature of this competition, especially on the point that IBM formed an organization designed to link knowledge of customers' business needs to knowledge of what was technically feasible in computing.

[24] This articulation of IBM's success draws heavily on Usselman (1993) and on my work with Franco Malerba (1998).

of a narrow definition like the speed of the computer. More important, the structures created by IBM to feedback user needs into technical improvements—to create new entrepreneurial knowledge—led to many new product features and technologies.[25] This planned initiative succeeded admirably until the late 1980s. Even the difficult transition out of the IBM mainframe computer era into the current "server" era was characterized by scarce entrepreneurial knowledge. I do not treat that transition in detail here, though Shane Greenstein and I have argued (1996) that its information needs were daunting and that the relevant information was highly dispersed.[26]

13.5.1 A (Different) Example Where Entrepreneurial Knowledge Was Less Scarce

It is worth pointing out the contrast to another GPT cluster in the computing industry that did not supply business data processing customers, but instead supplied technical, scientific, and engineering customers. The "minicomputer" industry was staffed by scientists and engineers and its customers were also primarily technical people, with technical problems to solve. Thus the minicomputer industry followed reasonably directly out of the original scientific and engineering knowledge basis of the electronic computer. Based on technical people selling to technical people, the minicomputer industry did not need elaborate structures to create new entrepreneurial knowledge. The relative scarcity and importance of entrepreneurial knowledge in WCA explains much of the difference of firm and industry structure between the business data processing sector dominated by IBM and the much more competitive minicomputer segment.[27] Ironically, the same shortages of entrepreneurial knowledge about customer needs that made scientific openness essential to the invention of business data processing (BDP) made entry and competition against IBM's position, once established, very difficult.

If not for the recombination of the electronic computer into a business data processing machine, society would "only" have gotten the kinds of

---

[25] Perhaps the most important solution to the problem of scarce knowledge about applications/technology overlap was IBM's invention of the closed, modular platform. This invention reduced the risk of customer experimentation dramatically. If a customer discovered that a particular business application worked, but that it required a larger or smaller computer, larger or smaller data storage, and so forth, they could move to those components without losing their initial investment in invention. This supported one of the most important forms of experimentation in business data processing, the construction of a complex high value system on top of a simple system. A customer might build an accounts receivable system that just kept track of who owed what, and then build a complex decision-support system on top of it to guide the extension of trade credit. If the trade credit system worked, IBM could offer the larger computers and data storage, and so forth needed to run it in a modular fashion.

[26] Bresnahan and Greenstein (1996) concluded from our empirical analysis that the most valuable computer applications were also the most difficult to invent given a new computer technology. We also concluded that technical progress in computing and technical progress in the uses of computing are very different bodies of knowledge.

[27] I am grateful to Shane Greenstein and to Franco Malerba in this regard; without our collaborations I would never have come to understand this.

$V(A,G)$ returns we got from scientific and engineering computing mostly supplied with minicomputers, not the much larger value associated with BDP. At this stage it is perhaps useful to reiterate what $V(A, G)$ means in this chapter. As a first point, what is *not* important is a judgment about the ultimate social value of business data processing versus scientific calculation. Instead, it is the area under the demand curve for BDP versus scientific, engineering, and other technical calculations (which takes the budget for science as given). Whatever the ultimate importance of science, science had significantly less willingness to pay for computers than did commerce over the second half of the twentieth century. However, at the crucial moment when the computer was being invented, scientists and engineers had the entrepreneurial knowledge (and the military demand) to fund the invention.

13.5.2 Invention of the PC As a Business Tool

The personal computer has found new bodies of demand a number of times. I focus here on the circuitous route to the first large markets for the PC as an individual productivity tool for white-collar workers.[28] As with other GPTs for WCA, it followed a circuitous route.

I revisit the familiar history of the very early PC industry with analytical goals in mind, taking repeated advantage of the gap between what we now know about the uses of the personal computer and what industry participants knew during the 8-bit era, roughly the late 1970s. That lets us understand the role of the information structure of invention at the time. The critical event still in the future was the invention and widespread distribution of personal productivity applications for white-collar workers. Market events during the 8-bit era were based on contemporary knowledge of demand—and on contemporary uncertainty about the future of demand.

That information structure of invention helps explain a number of market outcomes in the 8-bit era. Those include the importance of entrepreneurship, market selection of the more open platforms, firms' motivations for supplying open systems, and recombination. Accordingly, I will start with investigation of contemporary information, and then turn to examination of the supply of the two most successful platforms of the era.

There is real analytical value in understanding what suppliers did not know in the early days of the industry. That lets us understand firm strategies which were enabling rather than a planned initiative. It was commonplace in the 8-bit era to think of the main market of the PC as being hobbyists. Here is Microsoft founder Paul Allen in 1977: "The personal computer user finds himself at the leading edge of a new computer applications and technology, He is becoming a source of expertise and innovation. He is

---

[28] The history of these advances is carefully treated in a number of secondary sources, on which I rely heavily in this section. My account draws on Freiberger and Swaine (2000), on Ceruzzi (1998), Aspray and Campbell-Kelly (2004), on Langlois and Robertson (1992), and on other secondary and contemporary sources.

not merely a passive, casual user of hardware and software developed by others."[29] Around the same time, the founder of the commercial PC industry, Ed Roberts, forecast most business growth in "inventory, accounting, that sort of thing" (i.e., IBM mainframe-like applications for small business). With the candor and self-confidence characteristic of important inventors in computing, Roberts pointed out that no one present at the founding had a solid forecast of later developments (most pointedly, not his collaborator Bill Gates).

The most important platforms of the 8-bit era, commercially, were the Apple II and CP/M computers (running the CP/M operating system on a wide variety of brands of computers). Apple had a sponsored platform but a very open approach to developers. The design of the Apple II made it a mass market PC. The computer came in a plastic case, not metal, and looked like an office appliance more than a hobbyist's technology. It required no soldering, had a keyboard and a monitor, and could run programs. As a result the Apple II was dramatically easier to use than earlier personal computers (though still quite difficult to use by modern standards). Accordingly, it appealed to a far larger market than the hobbyist kits could. An important differentiator for Apple was that it used color, which appealed to game developers, but it appealed to the home and school user as well. On the other hand, the Apple II had a 40-column screen, fine for games and school but very problematic for word processing and spreadsheets. These design trade-offs reflected current technical levels, of course, but—as would be realized later during a scramble to make different trade-offs—also the key fact that demand forecasts were for hobbyist, home, and game.

Ken Olsen, founder and chairman of Digital Equipment Corporation (DEC), famously said in 1977, "There is no reason anyone would want to have a computer in their home." This remark is universally quoted to show that Olsen missed the opportunity represented by the PC. That dinosaur! This gives us an opportunity to be clear on who foresaw what. Contrast with Olsen's remark a contemporary explanation from Apple computer about the uses of its new PC, in a press release:

> Applications include using the computer as a teaching aid for students and for entertainment through interactive games . . . paddles and joysticks can be interfaced . . . a built in speaker sounds when the ball is hit or a photon torpedo is fired at a klingon. Manufacturers [Apple] also suggest home business applications such as financial and bookkeeping analysis, charting the Dow Jones average and home budget tracking. . . . [W]hen the Apple II is equipped with soon to be announced added components, it will be able to monitor home systems such as heating, cooling, burglar alarm, fire and smoke detectors and lighting. When you're away,

---

[29] "Software Column" by Paul Allen, VP of Microsoft: Personal Computing, January/February 1977, 66. At the time, Allen was Microsoft's "big think" person, while Bill Gates was more in charge of implementation.

> the computer can randomly light parts of the house on different days to give the appearance that someone is in residence.

Apple's description of the uses of its machine in this quotation include (a) immediately visible uses (games and educational software); (b) uses that still have not had any widespread commercial importance for the PC (burglar alarms, home heating, lighting, and cooling); and (c) uses that would find a mass market a decade or two later (home finance, which would become a mass market after the introduction of Quicken, and mass market use of online financial services, which would come with widespread use of the Internet).

The other main platform sponsor, selling CP/M, did not have Apple's marketing savvy, and simply admitted that it was up to others to figure out what the PC was for. "Statistics" and "Economics Research" were among the top uses of CP/M machines in a survey, suggesting a market somewhat smaller than 100s of millions of PCs. The point is, it was not merely Apple and DEC who lacked what we now know was key entrepreneurial knowledge about the use of PCs in offices. The lack was universal.

The founders of the PC industry did not particularly have white-collar automation in mind. (Except in the sense that they had everything in mind.) The first important platform sponsors in the PC industry, who built substantial (hundreds of thousands of units) commercial markets did not particularly have white-collar automation in mind. This leads me to the second central point, the widespread distribution of knowledge.

It was the invention of the word processor and the spreadsheet by new inventors—not the founders of the industry, nor people they had ever met—that turned the PC toward WCA. Interestingly, even the first inventor of a PC word processor, Michael Shrayer, who wrote Electric Pencil as a tool for printing manuals for his *real* software products, developer tools, did not really have WCA in mind. He had the immediate need to print manuals.

However, the creation of the PC and of a nonkit PC (Roberts, Jobs, and Wozniak at Apple) and of key software (Gary Kildall at Digital Research Inc., Gates and Allen at Microsoft) led to the creation of an enormous amount of market K. This, together with the open systems approach of early PCs, led to an explosion of applications software, but most particularly to the invention and commercialization of software for WCA. The inventor of the first spreadsheet, VisiCalc, absolutely had the automation of accounting work in mind. So did the effective commercializer of word processing, Seymour Rubenstein, seller of WordStar, who quickly entered and competed away Electric Pencil's business. The invention and commercialization of these very widely used applications turned the PC into a tool for the individual white-collar worker in the corporation. They were not anticipated by the founders of the industry. Indeed, once the inevitable consequences of the conversion of the PC into a white-collar tool occurred—IBM's entry, the professionalization of hardware and software supply—many of the founders reacted very negatively. Far from planned, this was a market outcome. If I have mentioned many of the inventors, it is to drive home the point that knowledge was very distributed and that decentralization was essential.

The entrepreneurs of WordStar and VisiCalc built large volume (by then PC standards) businesses because the main PC types, the Apple II and CP/M machines, were open to it and had rapidly growing installed bases. Existing PC firms—neither the inventors of the Apple or of CP/M, nor Micro Instrumentation and Telemetry Systems nor Microsoft, themselves pioneering and entrepreneurial—did not invent the new markets, nor did they commercialize them. The shortcomings of these firms (and of established firms like IBM and DEC) were not a limitation on what the market system could accomplish, however, as new firms opened up the new markets. Existing personal computer industry firms were a source of trained managers and potential distribution partners and technical collaborators for the new firms. This specialized and loosely linked structure worked well. It did not need planning nor central coordination to gain economies of scale in multiple products.

Through this inversion, a very valuable GPT cluster, the PC industry used (primarily) by white-collar workers was invented. Once again the first inventions served a technically-oriented community, hobbyists and hackers, with narrow goals. This time, that community was not academic science or military demand, but a self-organizing group much like modern open-source movement. They used some of the organizing principles of open science, however, including open systems. Some entrepreneurs would have liked to close systems, but the resource constraints of small firms in a small market left them compelled—recognizing that they did not know everything—to let outsiders innovate. Not only was there a shortage of entrepreneurial knowledge, the shortage was recognized and impacted business practice in a first order way.

With an important overlap between technical possibility and demand needs seen by no one, the early PC industry followed a circuitous route. The original invention for hobbyists, and the commercialization for home users were inverted by the invention of the word processor and the spreadsheet. This invention was inherently decentralized, as early movers did not anticipate what followed, and it led to a profound acceleration once the high value business PC markets were identified.

13.6 Major Mass Market E-Commerce, E-Content, E-Communication Initiatives

I turn now to the invention of a successful mass market platform: electronic commerce, content, and communications (hereafter EC^3), the widely-used Internet. To date, the Internet is the most important technology for the extension of WCA into markets. This famous example of recombination—the Internet had been used for other purposes for decades—lets us address two important analytical areas.

First, examining this invention, and the many failed planned initiatives that proceeded it, permits us to sharpen the concept and role of entrepreneurial knowledge considerably. The list of failed planned initiatives is remarkable, both remarkably long and remarkable for containing highly capable, knowledgeable firms with many resources. They all had *almost enough* entrepreneurial knowledge to start an EC^3 GPT cluster. They all knew that there was a broad mass market opportunity to create some kind of EC^3 GPT cluster. As in the case of the founding of the PC industry described earlier, we can take up the question of what inventors did not know when they did not know it. While all the planned initiatives failed, the actual creation of the successful EC^3 platform on the Internet was an inversion, following a highly decentralized, circuitous route. Attributing the success of the ultimately successful set of inventions to superior knowledge and foresight on the part of its early inventors is incorrect. Instead, the inversion was, as we shall see, a market work-around of important shortages of entrepreneurial knowledge.

Second, this important example lets us examine the role of openness in platform creation. This is considerably sharper than the cases examined before, because in this example entrepreneurial knowledge was not terrible, just not sufficient to permit success. Many of the planned initiatives were closed in ways that would have served the interests of platform sponsors or other early participants. Even when they were not extremely closed, and even when entrepreneurial knowledge was not terrible, they failed. The interaction between modest shortcomings of entrepreneurial knowledge and modest departures from open systems worked to block innovation. At the end of this section I discuss the theoretical salience of this finding.

The same history also shows that a circuitous route can invent something that is not obvious. Here I focus on two aspects of the widely used Internet. An innovation that satisfies a long-felt need, unsatisfied by many prior innovation attempts, is likely nonobvious. When the last key invention in the successful innovation is, from a strictly technical perspective, not a hard problem, the inference of nonobviousness is overwhelming. We shall see that the entrepreneurial knowledge needed to design a successful mass market EC^3 platform is what rendered it nonobvious. Open-systems innovation, which economizes on scarce entrepreneurial knowledge, was the key to success.

13.6.1 E-Commerce, Notably in Finance

The potential social value of mass-market electronic commerce was a long-felt need for many years before the widespread use of the Internet. Potential innovators knew that there was value in a platform for mass market electronic commerce. What they did not know, with adequate precision to guide a planned initiative, was the technical features of that platform and its relationship to other uses.

Mass-market e-commerce was a long-felt need in part because of the earlier success of e-commerce outside mass markets. Decades before the widespread use of the Internet, treasurers at large corporations could have access to bank account information electronically. Similarly, an airline reservations system could be accessed both by employee sales agents and by external (to the airline firm) travel agents. There were also some limited e-commerce applications that were used by the consumer, such as bank automatic teller machines. These applications crossed the boundary of the firm, which is why I call them e-commerce. What they did not do is reach a mass market of individuals using a common device. These applications did make it clear that one goal for WCA was crossing the boundary of the firm and automating markets (most white-collar work is in buying and selling bureaucracies). The invention and widespread adoption of the PC suggested to a wide variety of potential innovators that a GPT cluster of mass market e-commerce applications was feasible.

Many firms engaged in retail finance (banking and brokerage) saw this opportunity in the 1980s and first half of the 1990s and attempted to create a GPT cluster to fulfill it. These were not trivial undertakings, and often involved very large investments by very successful retail banking and brokerage firms, such as Chemical Bank, Bank of America, Banc One, Shawmut Bank, and so on. They also included Citibank, which had successfully pioneered the ATM network, one of the most successful mass-market e-commerce applications (but without a general-purpose "client" device) of the prior era. Many of these firms made very substantial investments in systems, and through much of this long era, these initiatives were always about to succeed. A 1983 article in *Time* entitled "Armchair Banking and Investing" (Alva, Ungeheuer, and Koepp 1983) pointed out that

> Bankers believe that financial services will eventually be part of futuristic home information packages like Viewtron that supply everything from recipes to movie reviews. Therefore they are scrambling to organize joint ventures with communications firms.

You can see from that very brief 1983 quote that the shock of the Internet was not the "vision" of delivering mass market EC^3 to consumers. These very early initiatives failed, as did their successors over the dozen or so years between this quote and the success of the Internet. One might think that the initiatives were technically too early or the attainable market too small before PCs diffused. However, over the relevant time period PCs got easier to use, diffused very widely, and became connected on better and better modems.

If the "vision" was present, what was the bottleneck? How was the bottleneck removed by the widespread use of the Internet? The *Time* quote, like many discussions by contemporary observers over the next dozen years, has several clues. Contemporary observers thought that a mass market financial e-Commerce system would need to be part of a larger "package" of online services to attract sufficiently many users to be economic. Bankers and brokers believed, rightly, that banking/finance applications alone, including checking brokerage account balances, online trading, online banking, and online bill paying, did not appear to offer enough value to end users.

The bankers and brokers solved this by turning potential collaborators in "home information systems" to offer users a "package." Conceptualizing the offering as a "package" for consumers captures much of the thinking at the time; that is, a planned initiative led by a consortium of applications developers. Turning to information firms for "home information systems" brought more knowledge of demand into the planned initiatives, a topic to which I now turn.[30]

13.6.2 Electronic Content (Mass) Markets

The potential social value of mass-market access to information and entertainment online was also, as a broad general idea, obvious for many years. There had been a number of online information systems in smaller markets, and their diffusion to mass markets was broadly forecast. There were even platforms for the sale of specific information services to their markets, and the expectation that a similar platform would emerge in the mass market arena was widespread.[31] None took off. This, too, permits a deep investigation of what the many failed potential innovators knew, and did not know, beforehand.

The conceptualization of many initiators of home information systems closely followed that of already existing business information systems of offering a subscription "package." High value information that already exists somewhere (stock prices on trades in the last 20 seconds) was already being sold at high prices to specialized audiences (traders, by Bloomberg). Surely lower value information that already exists somewhere could be sold to a mass market. For example, the editorial content of *Readers' Digest*

---

[30] This section has emphasized a mass market platform for home use because of the dramatic growth in home use post-Internet. There were, however, parallel initiatives for at-work use, also of limited success pre-Internet.

[31] A number of special-purpose online services had prospered, selling high-value information in narrow markets. One thinks of Lexis/Nexis selling information to attorneys, Bloomberg to the financial industry, DIALOG, and so on. By the late 1980s, there were hundreds of online databases. DIALOG was a database platform; searchers and readers would pay between $35/hour and $500/hour depending on the database. Bloomberg, founded in 1981, was founded by a former financial market participant (at Salomon brothers) who saw the benefits of delivering already existing information to financial market participants. They would lease a "Bloomberg machine" (i.e., a special-purpose terminal), and get rapid 24 hour access to financial and related information. These successful commercial online services had themselves been invented by circuitous paths (e.g., DIALOG started at Lockheed).

already existed in machine readable form: surely it could also be sold somehow at lower prices to a mass market online audience? The *Readers' Digest* example is real, and a large number of publishers of consumer-oriented media content sought to move online over the 1980s and early 1990s.

Many of these existing publishers of consumer-oriented media recognized the limitations of their entrepreneurial knowledge and sought to overcome those limitations by undertaking joint ventures or alliances with technology firms. Knight-Ridder, CBS, and Times-Mirror all had collaborations with AT&T. Many other firms had collaborations with IBM. Harrigan (2003) has a very useful review of the wide list of joint ventures (JVs) and alliances that arose in this area. Like the other media firms that sought to create mass markets on a go-it-alone basis, these collaborations did not succeed in creating a mass market.

The plethora of attempts at mass-market e-content typically set up the online services as closed, with particular attention to the unauthorized copying of content, which often gave control rights to the owners of a particular kind of content. While those contractual protections may have had a good economic purpose looking only at local knowledge, they were problematic for creating a broad general GPT cluster involving different kinds of content and service. The other potential suppliers of e-commerce services, for example, would not necessarily have adopted a subscription model nor would they have emphasized the prevention of copying. Making this problem more difficult—as we now know from watching the struggles of "content" providers from magazines to Hollywood adapt to the Internet, the iPad, and so forth—is that the entrepreneurial knowledge of exactly how existing content will be sold in a new medium is hard to come by. How much harder when the medium has yet to be invented! There were many of these initiatives, spread out over a wide variety of content companies, joint ventures with existing telecommunications companies, and computer firms. I will not attempt a complete list here because the economically important point is that, even taken together, these initiatives did not attract sufficient end-user interest to start a positive feedback loop around mass-market e-content.

13.6.3 Electronic Communication for Mass Markets?

Similarly, a wide number of firms offered electronic communications services to consumers and/or to firms in the period preceding the widespread use of the Internet. Many of these looked like modern e-mail, and indeed shared some technology with the development of e-mail in not-for-profit settings on the Internet. None of the for-profit ones were as large as the user-built e-mail network serving existing Internet users (largely in universities and related places). The end result was also low usage, and the network effects of communications systems create much more value in widely used systems. By the early 1990s, one could see the odd result that scientists and engineers, surely not the most communicative of people, had excellent access to e-mail on the Internet, but that other classes of users, whether as employees or as consumers, had much more limited access. This makes it clear that mass market electronic communications was also a long-felt need. Direct efforts to push it to firms and consumers were, however, proceeding slowly.

I have reviewed just a few of the many planned commercial initiatives in the dozen or so years before mass use of the Internet took off. Many firms were throwing large R&D budgets at one aspect of EC^3. None of them had quite the right knowledge to pull it off; all knew the social return was high, but no one could find quite the right direction of technical progress to unleash it. In this long era, technologies that might make the PC into a communications, real-time entertainment, or information gathering tool existed but were narrowly distributed. The Internet ones were narrowly distributed to academic and related communities. The commercial ones were narrowly distributed because of their proprietary or top-down nature. There were huge network effects benefits that could follow from a data communications network—being able to e-mail pretty much anyone, for example. Yet these remained latent because no network was ubiquitous.

13.6.4 Planned Initiatives As a Coordination Device

The previous subsection pointed to a number of mass market EC^3 initiatives that were most strongly pushed by a particular kind of application. Bankers pushed mass market e-commerce, publishers pushed mass-market e-content, and technology firms pushed mass market e-communications—and many others not reviewed here. None drew a widespread enough audience to ignite a mass market. This problem of fragmentation was not lost on contemporary observers who noted that, to attract sufficiently many consumers to create a positive feedback loop, e-commerce sites would need e-content and e-communications services, and vice versa. We now know that this problem was solved by the Internet inversion, which drew in sufficiently many users to create many opportunities for all three kinds (C^3) of applications both reaching consumers and workers, and whose openness permitted rather than coordinated the supply or applications.

One might think that this problem could be solved by coordination and the creation of a general mass market online platform. The most important lesson of mass market EC^3 is that this intuition, too, is wrong when entrepreneurial information is scarce. To see this, I now examine the two most successful planned initiatives led by a platform sponsor before the widespread use of the Internet, America Online (AOL), and Microsoft Network (MSN).

Each of these was an “online service,” meaning a closed, proprietary platform for EC^3 applications. Online services provided infrastructure for EC^3 applications. They were set up to take advantage of central control of the platform. Following ideas like those in the “two-sided markets” literature, online services would have contracts both with applications developers and with users. They would collect revenues from the users and pay the developers. Control permits complex pricing schemes in such a platform. Users typically paid a monthly subscription fee, and could also pay by the minute they were connected to the service or value added charges based on what services they used, content they looked at, or applications that they ran. A large service could license in many applications from a wide variety of third-party inventors. Online services also provided infrastructure so that subscribers could communicate with one another. For example, they may have e-mail services or online discussion areas or forums. Each online service was a closed system, in competition with the other closed systems, and content was typically local to each online service (though there was some multihoming) and the communications services offered were also local to the specific online service.

While online services followed the program suggested by the "two-sided market" literature in economics—that is, a benevolent dictator platform sponsor offering complex prices to both sides (users and applications) and competing with other platform sponsors, they were only moderately successful. That is not to say they failed as businesses, but all of these online services now seem to us to be smaller, less rich, and more expensive than the commercial Internet.

The most successful online service for end consumers before the widespread use of the Internet was AOL. America Online was marketed to consumers as a general online service, and it provided e-mail (to other AOL users) and related communications services. America Online also offered content providers and e-commerce merchants the opportunity to put materials inside AOL's "walled garden." America Online would then distribute those materials online to consumers. Startup AOL was not the only online service, as computer heavyweight IBM and retailing heavyweight Sears collaborated to build one. Many firms saw the broad, general opportunity.

America Online was successful enough to draw competitive imitation from Microsoft. Microsoft created an AOL-imitation online service, called MSN, which followed the walled-garden model. There would be e-communications tools for users, and authoring tools for e-commerce and e-content providers who wanted to sign a contract with Microsoft to share revenue. An important advantage of Microsoft's plan was the widespread distribution of the MSN "client" software, which, starting with the release of Windows 95, would be distributed with new computers, an obvious mechanism to build a mass market. The idea of widespread distribution to consumers was also responsive to the biggest problems faced by existing EC^3 initiatives; that is, getting enough users to attract a wide variety of developers. Another reason to examine MSN is that, technologically, it was newer than the widespread use of the Internet. When Microsoft launched it the Internet inversion was already almost completed. Microsoft Network did not fail because it used earlier-vintage technology nor because it had no good plan for mass usage. It failed because the "Internet tidal wave" rolled over it.

We did not get to see the AOL-MSN competition that would have followed but for the widespread use of the Internet. Both were quickly competed into irrelevance by the Internet. The MSN was withdrawn (confusingly, there was a later Internet website with the same name from Microsoft) and AOL became a "gateway" to the Internet. Absent the widely used Internet, would the AOL-MSN competition have led to widespread EC^3 with as much innovativeness, breadth of uses, and usage? While it is always difficult to answer a historical counterfactual, at least two important considerations make it clear that the likely outcome would have been significantly slower to develop, less innovative, less flexible and changing, and smaller than today's Internet.

13.6.5 Why Planned Initiatives Failed

The last pre-Internet initiative also offers us an opportunity to hear the insider perspective from Bill Gates of Microsoft on the disadvantages of MSN versus the Internet (emphasis in original):[32]

> Subject: Internet as a business tool
> I know I am a broken record on this but I think our plans continue to underestimate the importance of an OPEN unified tools approach for the Internet. The demo I saw today when Windows 95 was showing its Internet capability was someone calling up the Fedex page on the Internet and typing in a package number and getting the status. Imagine how much work it would have been for Fedex to call us up and get that running on MSN and negotiate with us. Instead they just set it up. A very simple way to reach out to their customers. The continued enhancement of the browser standards is amazing to me. Now its security and 3d and tables—what will it be within the next several years? Intelligent controls, directory—everything we are trying to define as standards.

Gates makes two arguments here that are salient to our inquiry. First, he sees the advantages of the permissive nature of a new application development in an open environment (in his discussion of Fedex.) The attempt to keep control slows innovation by lowering λ_2. Second, he sees the open Internet as being as effective as a planned initiative in creating a "unified tools approach" and in "continued enhancement of . . . standards." This is analytically important because many advocates of planned GPT initiatives assert that planning will produce superior architectures. There are, of course, cases in which planned initiatives are better in that regard, as we saw

---

[32] This is an e-mail from Gates on April 6, 1995 to a number of senior Microsoft executives including those responsible for MSN. It was published as a result of the Antitrust case and is located in Government Exhibit 498. I cite it as Gates (1995).

in the IBM business data processing example earlier, but open decentralized market innovation can be very good for standards, as it was in the PC. It can offer an important competitive alternative.

Latent in Gates' remark is also a serious problem of centralized contracting. Solving fragmentation through a planned initiative would call for entrepreneurial knowledge of the possible developments in e-content, e-commerce, and e-communications to attract many complementary applications, and also for sufficient knowledge of the relevant consumer marketing issues to create a widespread mass market. That is a lot of entrepreneurial knowledge to get together in one place. Openness economizes on it—no one would need to know the potential invention possibilities at Fedex and at millions of other firms to know how to structure the platform contract.

A related point about the difference between walled gardens and open systems is the potential for transformative recombinant innovation by providers of complements. We saw this in the PC example and also here. The open Internet has given us a wide number of innovations that run on the server; one thinks immediately of Yahoo, Google, Ebay, Amazon, WikiPedia, and Craigslist. The first four of these would have been perceived as duplicative or as competitive threats by a walled garden online service provider, and the last two would have faced difficulty at the time of their founding, paying for space in a walled-garden environment. The distributed innovation essential to the acceleration of an inversion would have been problematic for MSN or AOL.

Another reason to believe the pre-Internet initiatives would have gone less far and much less fast is that their proponents anticipated a long, slow growth path. Microsoft, for example, thought that the diffusion of broadband connections to the home would be an important growth driver for MSN, and was (wisely, given their entrepreneurial knowledge) investing in online systems in advance of that development. Broadband diffusion would have been even slower than it has been historically if not for the explosion in telecommunications demand driven by the Internet.

13.6.6 The Internet Inversion

In the aforementioned, I noted many participants who lacked entrepreneurial knowledge at an early stage. It is worth considering how knowledge changed as a result of the Internet inversion.

To begin, let me very briefly recount the familiar steps leading to an Internet suitable for mass-market use. After beginning as a military technology, the Internet spent much of its youth as a partly National Science Foundation (NSF)-sponsored network in universities, military installations, and some technical companies. In this era, a number of important developments occurred, including valuable add-on facilities for e-mail, for discussion and "social" networking (like Usenet—which is "social" in the sense engineering communities can be, not in the sense of Facebook), and for sharing data sets and the like among scientists. Two important steps moved the Internet closer to mass market use. The first was the creation of the World Wide Web (WWW) in the computer department of CERN, a physics laboratory. The World Wide Web runs on top of the Internet and provides for a system of interlinked hypertext documents. The WWW was clearly envisioned by its inventors as entirely general (like a number of other networks of the era) and had several useful features that permitted generality, including the use of URLs, a broad open capacity for adding materials, and so on. The application that paid for the development of the WWW, however, was to permit physics researchers to share data sets.

The final step toward mass market use was the invention of the web browser at another computer department of another physics laboratory, this one at the University of Illinois. The web browser was almost purely a recombination of existing elements. However, to quote Schumpeter again, while there are "numerous possibilities for new combinations" they are only obvious ex post. Before the recombination, ex ante, "most do not see them." As a technical matter, the browser's inventors recombined the idea of a graphical user interface with some inventions and improvements in that interface (the "back" button) with existing hypertext protocols. This was the last step in the inversion that was entirely within the technical world, and it was adequately simple to invent that the resources available to one physics lab at one university could pay for it.

The web browser and the open WWW were sufficiently suitable to mass market that they began to draw many users, creating the so-called "Internet mania." A number of applications were quickly available, many involving user-generated content. The availability of e-mail as an already developed application—and a free one—was also a driver of rapid adoption.

One of the inventors of the browser first searched for jobs in interactive television, the Silicon Valley rage of the moment, and then became a founder of Netscape, the commercializers of the browser.[33] (Entrepreneurial knowledge is about overlaps, not about envisioning the whole thing.) A venture capitalist who backed Netscape, L. John Doerr noted the dramatic change in the state of knowledge after the creation of the noncommercial "Mosaic" browser:

> I'd seen Mosaic, the UNIX version of it. . . . Marc earned $3.65 an hour, or whatever the University of Illinois had paid him . . . and 2 million people were using it. You would have to be dumb as a door post not to realize that there's a business opportunity here.

---

[33] There were many, many false starts for online content. I have skipped the enormous category of them related to "convergence" of traditional mass media with computing. That a key figure in the commercial development of the Internet mass market almost worked in one of them is as telling about entrepreneurial knowledge as the broad ignorance of WCA possibilities at the founding of the PC industry.

That is the hallmark of a change in knowledge, ex post obviousness. Decentralization was essential here again, as the commercializers of the browser quickly drew criticism from Internet and especially WWW technologists (in much the same spirit that many inventors of the early PC or of early computers criticized IBM) for being commercial. The problem with open systems from a first-inventor perspective is not that recombination may create something unanticipated, but that it may create something undesired.

Mass market electronic commerce, content, and communication is one of the great triumphs of recombination. It represents a dramatic increase in the value-in-use of a wide number of preexisting technologies, from the telephone network to the PC, from the server and the database management system to the marketing knowledge of a number of existing retailers. The invention of those preexisting technologies was financed with knowledge of and in anticipation of their own original markets, not primarily in anticipation of mass market EC^3 returns, and their recombination represents a social boon.

Mass market EC^3 was triggered by a series of GPT component invention: the browser, the WWW, and the Internet. Each of these was invented or innovated in low-resource environments but environments where (a) entrepreneurial knowledge showed how a particular problem could be solved in a general way and (b) openness was a natural way to compensate for resource scarcity.

The Internet mass market platform for EC^3 has several important features that sharpen our understanding of entrepreneurial knowledge in the case of "platform" industries. The failures that preceded the mass market use of the Internet had the feature that many firms

- Knew that some kind of platform for mass market EC^3 would be valuable.
- Knew that any such platform would need to recombine some aspects of business data processing, telecommunications, and the PC.
- Knew many of the applications areas in which value would arise.
- Did not know, however, what mix of applications (i.e., what services, content, and e-markets), would draw mass user participation.
- Did not know what "business models" would be successful in many of the key applications sectors.

The problem of entrepreneurial knowledge is knowing what product will sell in a new market. This problem is ratcheted up in the platform creation or GPT context. A platform entrepreneur needs to know what group of applications (including content) will attract a group of users that will in turn be attractive to creators of the relevant applications. The scope of knowledge required ex ante appears daunting. Small surprise, then, that there have been elements of decentralized exploration, even to the point of inversion, in the creation of many important platforms. The opportunity to recombine, as much as the vision to create, are central to the invention of many of the most valuable modern GPT clusters.

13.6.7 Relationship to Recent Literature

If the planned initiatives followed the directions of the "two-sided markets" literature and failed for that reason, where is the gap? The literature investigates the benefits of creation of a platform sponsor creating a set of incentives for market participants, under the assumption that the appropriate platform sponsor is known, or can be determined by ex ante competition, and has sufficient entrepreneurial information (which need not be perfect) to set incentives for participants, including incentive to invent applications. Thus, for example, Weyl and Tirole (2010) have a careful treatment of the relationship between the social optimum incentive scheme and the one that would be picked by a platform sponsor. These point out that an effectively designed incentive scheme can effectively reward applications or content creators, and that the platform sponsor is in a position to create and to benefit from an incentive scheme that benefits both users and creators. Like earlier work by, for example, Baumol and Willig (1981), they note that incentives to discriminate across groups can be quite good for discriminating monopolists. Their central policy proposal, creation of a local set of incentives by a platform sponsor, is also the best explanation of how the mass market EC^3 planned initiatives studied in this section failed.

The important point is that innovation sometimes calls for decentralization, not planning. The path to creating a new platform often calls for shifts in leadership, something that cannot be left to a platform sponsor as their incentive is to maintain leadership. The creation of new platforms, under conditions of distributed knowledge, calls for permitting not coordinating. Both of these economic effects take us outside the assumptions of the "two-sided markets" literature.

The history of efforts to start mass-market electronic commerce, content, and communication is revealing about the knowledge needed for a planned effort to create a new GPT cluster. The first successful mass-market e-commerce, e-content, and e-communication GPT cluster, the widely used Internet, emerged by a circuitous route marked by inversion. A long series of planned efforts to create such a GPT cluster failed. The planned efforts reviewed in this section were closed commercial initiatives that drew on the entrepreneurial and technical knowledge of some very impressive market participants. The failures, as we see in this section, arose because their entrepreneurial knowledge was limited, even though it was almost right. Examining them permits us to sharpen the concept of entrepreneurial knowledge considerably. It also shows, once again, the importance of openness in permitting multiple innovators to create what no single planner could.

13.7 Smartphones

Sometimes an entrepreneur has sufficient entrepreneurial knowledge for a planned initiative, particularly when many of the market pieces are already in place. West (2009), in an interesting history of the invention of the Apple iPod, shows that one firm succeeded in creating a digital music player with Internet distribution. Previous efforts had either failed with consumers or with music studios. Apple's understanding of the incentives of music studios was deep, and it takes nothing away from their accomplishment to say that songs, while hard to invent, have as a group more easily forecast demand than do WCA productivity apps. In this instance, the problem of entrepreneurial knowledge formation was solvable, and solved. Apple's formidable ability to design for consumer use, and canny observation that computer power and storage were now low enough for a device, were congruent with the needs of building a mass music platform. All the more impressive is the same firm's building upon that base to create a smartphone applications platform by building on the base in music and on the technical infrastructure put in place by mobile carriers. The mobile carriers had "app stores," but never ones with much volume. Important early applications such as games were, once again, not impossible to foresee, but one firm did see the platform opportunity with enough clarity about complementors' incentives to start a planned initiatives. That there are so few examples of successful planned initiatives illustrates the difficulty of coming up with sufficient entrepreneurial knowledge in computing in general.

In computing, the biggest shortages of entrepreneurial knowledge have arisen at the founding of the WCA GPT clusters, as we have seen. Founding science and engineering platforms or consumer entertainment platforms has been easier. The variation arises, not in the technology itself, but in the market problem of foreseeing what technology will bridge to the very hard-to-forecast automation of white-collar work in bureaucracies and markets.

13.8 Conclusion

The GPTs call for invention both in general components and in applications sectors. This raises the possibility that the founding of GPT clusters may, like recombination, be held back by scarcity of entrepreneurial knowledge. Ex ante, there may be no single locus of knowledge of the precise direction of technical progress into the overlaps between technical opportunity and growth needs. This lack of anticipation does not follow from irrationality or similar phenomenon, but instead reflects the distribution of knowledge across many agents in a market economy. Some know technical opportunity; others know the growth needs.

I have brought forth both a very simple theory of this and undertaken historical investigations to foreground an important fact about late twentieth century and early twenty-first century economic growth. The ex ante problem of scarce entrepreneurial information has led each of the major white-collar automation technologies in computing to be invented by a circuitous route of inversion, decentralization, and acceleration. Important recombinations of these technologies into new, more complex systems have also been characterized by much better knowledge ex post than ex ante. Since WCA will continue to be one of the central growth poles of the twenty-first century, this is an important lesson. Little can be done to solve the problem of scarce entrepreneurial knowledge in this area.[34] Much can be done, however, to preserve the openness and decentralization that have been so important.

Many observers are tempted to conclude that the Internet inversion, the general purpose computer inversion, or the PC inversion involved pivotal steps. To take the largest of three very large literatures, a number of observers have argued that the "countercultural" (in the 1960s political sense) communities involved in the development of the PC were pivotal. I admire the achievements of many countercultural inventors of the PC revolution, just as I admire the achievements of scientists in creating the computer or the widely used Internet. But we should be careful before we conclude anything was pivotal. The logic of an inversion does not say that the particular circuitous route taken to found any particular GPT cluster is pivotal. It is close to saying the opposite—there are a wide variety of paths to collective discovery of a valuable GPT. The "countercultural" nature of some PC innovators, the technical nature of many others, the military and scientific nature of key inventions of the general purpose computer (or Internet) innovators play two roles in the analysis. The first is that they are examples of diversity, especially diversity in entrepreneurial knowledge. The importance of diversity means that few are pivotal. Second, they used open approaches, often because of the very limitations of their entrepreneurial knowledge or their capabilities. Openness is crucial but likely no inventor was pivotal.

A similar problem applies to the common argument that small historical accidents in the founding of GPTs and in recombination are determinative of events for decades if not centuries afterward. While there was clearly some inertia around the IBM computer standard and there is some inertia around the Windows PC, those came at the exploitation stage, not at the earliest stages of exploration. More broadly, a decentralized and diverse economy will find and exploit large overlaps between technical opportunity and growth needs. The lesson we should take away from the particular paths used historically are first, that openness was important to market solutions,

---

[34] There have been numerous failed efforts over the last fifty years to improve ex ante knowledge about WCA. Most have used an engineering approach to organizational design or customer relations.

and second, that the apparent maturation of some industries (such as the IBM mainframe and, one can only hope, the Windows PC) can itself be an intermediate stage. Abandoning openness at this stage would be a major error.

References

Acemoglu, D. 2002. "Directed Technical Change." *The Review of Economic Studies* 69:781–809.

Aghion, P., and P. Howitt. 1998. "On the Macroeconomic Effects of Major Technological Change." *Annales d'Économie et de Statistique* 49/50:53–75.

Alva, Marilyn, Frederick Ungeheuer, and Stephen Koepp. 1983. "Armchair Banking and Investing." *Time,* November 14.

Arora, A., A. Fosfuri, and A. Gambardella. 2001. *Markets for Technology: The Economics of Innovation and Corporate Strategy.* Cambridge, MA: MIT Press.

Aspray, William, and Martin Campbell-Kelly. 2004. *Computer: A History of the Information Machine,* 2nd ed. Boulder, CO: Westview Press.

Athey, S., and S. Stern. 2002. "The Impact of Information Technology on Emergency Health Care Outcomes." *The RAND Journal of Economics* 33 (3): 399–432.

Atkeson, A., and P. J. Kehoe. 2007. "Modeling the Transition to a New Economy: Lessons from Two Technological Revolutions." *American Economic Review* 97 (1): 64–88.

Basu, Susanto, and John G. Fernald. 2008. "Information and Communications Technology as a General Purpose Technology: Evidence from U.S. Industry Data." *Federal Reserve Bank of San Francisco Economic Review 2008:*1–15.

Baumol, William J., and Robert D. Willig. 1981. "Fixed Costs, Sunk Costs, Entry Barriers and Sustainability of Monopoly." *Quarterly Journal of Economics* 96 (3): 405–31.

Bresnahan, T. F. 2010. "General Purpose Technologies." In *Handbook of the Economics of Innovation,* edited by Bronwyn Hall and Nathan Rosenberg, 761–92. Amsterdam: North Holland Elsevier.

Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *The Quarterly Journal of Economics* 117 (1): 339–76.

Bresnahan, T. F., and S. Greenstein. 1996. "Technical Progress and Co-Invention in Computing and in the Uses of Computers." *Brookings Papers on Economic Activity: Microeconomics* 1996:1–83.

Bresnahan, T. F., and F. Malerba. 1998. "Industrial Dynamics and the Evolution of Firms' and Nations' Competitive Capabilities in the World Computer Industry." In *The Sources of Industrial Leadership,* edited by D. Mowery and R. Nelson, 79–132. Cambridge: Cambridge University Press.

Bresnahan, T. F., and M. Trajtenberg. 1995. "General Purpose Technologies: Engines of Growth?" *Journal of Econometrics* 65 (1): 83.

Brynjolfsson, E., and L. M. Hitt. 2000. "Beyond Computation: Information Technology, Organizational Transformation and Business Performance." *The Journal of Economic Perspectives* 14 (4): 23–48.

Ceruzzi, Paul E. 1998. *A History of Modern Computing.* Cambridge, MA: MIT Press.

Cohen, W., and D. Levinthal. 1990. "Absorptive Capacity: A New Perspective on Learning and Innovation." *Administrative Science Quarterly* 35:128–52.

Dosi, G. 1982. "Technological Paradigms and Technological Trajectories: A Suggested Interpretation of the Determinants and Directions of Technical Change." *Research Policy* 11 (3): 147–62.

Fleming, Lee. 2001. "Recombinant Uncertainty in Technological Search." *Management Science* 47 (1): 117–32.

Freiberger, Paul, and Michael Swaine. 2000. *Fire in the Valley: The Making of the Personal Computer,* 2nd ed. New York: McGraw-Hill.

Harrigan, K. R. 2003. *Joint Ventures, Alliances, and Corporate Strategy.* Lexington: Beard Books.

Hayek, F. A. 1945. "The Use of Knowledge in Society." *American Economic Review* 35 (4): 519–30.

Helpman, E., and M. Trajtenberg. 1998. *A Time to Sow and a Time to Reap: Growth Based on General Purpose Technologies.* Cambridge, MA: MIT Press.

Jones, Benjamin. 2009. "The Burden of Knowledge and the 'Death of the Renaissance Man': Is Innovation Getting Harder?" *Review of Economic Studies* 76 (1): 283–317.

Jovanovic, B., and P. Rousseau. 2005. "General Purpose Technologies." In *Handbook of Economic Growth,* Volume 1B, edited by Philippe Aghion and Steven N. Durlauf, 1181–222. Amsterdam: Elsevier B.V.

Langlois, R. N., and P. L. Robertson. 1992. "Networks and Innovation in a Modular System: Lessons from the Microcomputer and Stereo Component Industries." *Research Policy* 21:297–313.

March, J. 1991. "Exploration and Exploitation in Organizational Learning." *Organization Science* 2 (1): 71–87.

Mokyr, J. 2002. *The Gifts of Athena: Historical Origins of the Knowledge Economy.* Princeton, NJ: Princeton University Press.

Nelson, R., and S. Winter. 1982. *An Evolutionary Theory of Economic Change.* Cambridge: Belknap Press.

Romer, P. M. 1987. "Growth Based on Increasing Returns Due to Specialization." *The American Economic Review* 77 (2): 56–62.

Rosenberg, Nathan. 1963. "Technological Change in the Machine Tool Industry, 1840–1910." *Journal of Economic History* 23 (4): 414–43.

———. 1996. "Uncertainty and Technological Change." In *The Mosaic of Economic Growth,* edited by R. Landau, R. Taylor, and G. Wright, 334–52. Stanford: Stanford University Press.

Schumpeter, J. 1939. *Business Cycles.* New York: McGraw-Hill Book Company, Inc.

Scotchmer, S. 2004. *Innovation and Incentives.* Cambridge, MA: MIT Press.

Usselman, S. 1993. "IBM and Its Imitators: Organizational Capabilities and the Emergence of the International Computer Industry." *Business and Economic History* 22:1–35.

Weitzman, M. L. 1998. "Recombinant Growth." *Quarterly Journal of Economics* 113 (2): 331–60.

Weyl, E. Glen, and Jean Tirole. 2010. "Materialistic Genius and Market Power: Uncovering the Best Innovations." Working Paper.

Comment Benjamin Jones

This chapter has a "big think" orientation and reveals numerous insights about the innovation process. The starting point is to recognize that knowledge required for successful innovation is distributed across many agents. These agents do not know each other, so their individual knowledge is not easily aggregated or shared. The chapter then makes a distinction between two types of knowledge that are relevant to innovation. There is technical knowledge—the actual engineering and scientific know-how to actually make something. And there is entrepreneurial knowledge—knowledge about whether there is a market for the new thing and, if there is a market for it, whether there might be other, associated markets and complementary recombinant innovations that further justify going down the initial path.

Note immediately that there are some standard innovation flavors here. There is uncertainty about the innovative possibilities. Ex ante, it is ambiguous what these innovative opportunities are, technically and in the market. There is also an emphasis on complementarity, both the interdependence of knowledge and the consequent interdependence of agents across whom the knowledge is divided up.

But the key addition of the chapter is a flavor of Hayek, asking how distributed knowledge can be brought together and emphasizing the role of the market. If someone actually delivers an innovation to the marketplace, then distributed agents see the innovation and recombine it with their own ideas. Before the innovation is delivered to the market, there is an absence of knowledge. The core idea in this chapter is that in making the thing and bringing it to the market, the information burden on everybody else is relieved. This action turns one agent's entrepreneurial knowledge—perception of a particular opportunity—into widespread knowledge, making it easier to recombine and build into additional innovations. Another theme of the chapter is that this process may be especially critical for general purpose technologies.

The following simple formalization can capture many of the main ideas in the chapter and demonstrate the generality of applications that emerge from Bresnahan's analysis. Imagine you are considering an innovation A with value $V(A)$, which can be obtained for a cost r. Furthermore, imagine there is some possibility of combining innovation A with some other innovation B, giving your initial effort some additional value $V(A, B)$. The analysis of

---

Benjamin Jones is associate professor of management and strategy at the Kellogg School of Management, Northwestern University, and is on leave from the National Bureau of Economic Research.

the chapter hinges on whether you can expect, in this world of decentralized knowledge, to capture this $V(A, B)$.

Write the expected return on the innovation A as[1]

$$\{V(A) - r\} + V(A,B) * \lambda * K.$$

Let your bargaining power be measured by $\lambda \in [0,1]$, defining the share of the additional income $V(A, B)$ that you would capture. Define $K \in [0,1]$ to represent the probability that you will perceive the opportunity of $V(A,B)$. The issue is thus partly one of bargaining power over future innovations, for example, due to intellectual property rights. The issue is also one of knowledge—you may not know that the recombinant possibility even exists.

The interesting case, of course, is where $V(A)$ is less than r. Then, on your own, you may choose not to produce innovation A given its expense. Yet there may be substantial value in the recombination of A and B. The challenge is either that you do not look forward to a large share of the value (λ is low) or you do not readily perceive the combination itself (K is low).

The bargaining problem suggests that you need high λ to encourage the investment in A. However, while high λ means that you can appropriate most of the market—the $V(A,B)$—it also implies that other would-be innovators become less inclined to create B, because now they cannot get much of the recombination benefit for themselves. This trade-off, and its implications, has been studied extensively by Suzanne Scotchmer.

The emphasis and novelty of this chapter surrounds the question of knowledge itself, represented by K. Even if we solve the bargaining problem, you still will not get innovation if K is low. The innovator has little or no idea what this B is. This lack of knowledge could surround technical aspects of B, market knowledge for B, and/or B's recombinant prospects with A. These possibilities may be very hard to foresee, especially when knowledge is distributed.

Returning to the Hayek theme, one (imperfect) solution to this knowledge problem is for someone to simply create B and bring it to the market. Then people see it, resolving the K problem, and now may create A. The marketplace thus helps unleash recombinant innovation.

The general purpose technology (GPT) version of this analysis is to imagine that there are lots of potential innovations that could recombine with A (the GPT),

$$\{V(A) - r\} + \sum_i V(A,B_i) * \lambda_i * K_i.$$

This setup suggests a natural story for "inversion" as the initial step for the spread of a GPT. The GPT is originally produced with a narrow application

---

[1] This notation and setup is not quite what was used in initial drafts of the chapter, but is simple and sufficient to capture some key ideas.

in mind. This is the case where $V(A) > r$ and the innovation goes ahead without consideration of the recombinant possibilities. For example, computers were originally developed to perform narrowly defined calculations, and government researchers created the precursor of the Internet for their own narrow purposes. There was little knowledge about the ultimate potential (the K_i were low). But having produced A, these areas started witnessing decentralized innovation. While the A people did not see the B_i—and likely were not even thinking about B_i—suddenly there are all these agents thinking about A, because now they can see it. So the decentralized B_i people dive in and innovation accelerates; if A is a general purpose technology, decentralized innovations can really take off.

The rest of my comments will depart from the general purpose technology focus of the chapter and consider some other applications of this simple framework, which can further demonstrate its use.

Consider basic research. Basic research typically shows little or no market value directly ($V(A) < r$) but may have lots of recombinant possibilities for commercial innovations (the $V(A,B_i)$ may be large). That is often how economists describe basic research and the reason it may be underprovided. The standard policy solution is subsidization: public institutions pay scientists a wage and provide research funds. In addition, we make A freely available: we set $\lambda = 0$ for producers of basic research. Thus the distributed Bs capture the full value of recombination, incentivizing their activity. This perspective provides a standard description of the "public, open science" model, which is a good description of many national innovation systems.

The additional nuance that Bresnahan's approach reveals centers on the dissemination of basic scientific knowledge. With basic science, the output is not presented as a standard good or service, demonstrating revenues, costs, and profits in the marketplace. Rather the output is a paper, a seminar, an informal chat with colleagues. How does the commercial market learn about the new idea or whether it is valuable? That is, how successfully does the "public, open science" model solve the "low K" problem? Papers and conferences are part of the solution but may be incomplete; for example, they do not convey tacit knowledge. One solution for commercial enterprises may be geographic agglomeration around universities. Private firms locate around Stanford and Berkley, MIT, and so forth, explicitly to increase their K.

In this view, the effectiveness of agglomeration will depend on the capacity of private firms to search the university for good ideas. That is, the agglomeration solution—using a local network in place of an arms-length market—is not the Hayek-like solution. Recall the starting point of the chapter—knowledge is distributed across agents. The market may solve this problem when an innovation is sold, but if direct communication is important to acquire basic science ideas (hence explaining agglomeration) then firms' acquisition of researchers' ideas depends on the researcher's willingness to engage. If the researcher's interests or incentives are defined by producing additional basic research, why exactly does the researcher take the firms' phone calls? Does the researcher want to spend hours and hours talking to private firms?

Here the issue of openness becomes more complicated. Namely, the K issues and λ issues start to interact. Can you tell the basic researcher "you have to publish your ideas for free" ($\lambda = 0$) and also say "you still need to take calls from all these commercial people who are going to make all the income from your idea"? That is not easy. So perhaps we need to think about giving some λ back to the basic researchers, which would result in higher K for others.[2] Alternatively, we can imagine that firms will simply pay researchers for their time (i.e., a consulting fee), which would also raise K. This solution might be difficult in practice, however, given the substantial compensation and costly bargaining that might be needed with each researcher, the breadth of search the firm must undertake, and the bias expansive (and thus expensive) search may impose against small firms. These are key questions for understanding possible market failures in the commercialization of university research and the ultimate returns to basic science.

One can also think about standard setting through this lens. Think of a standard as an innovative output, A. By publicly agreeing to A, the market enhances recombinant possibilities by raising K. This knowledge is not standard marketplace knowledge based on profits from a new innovation, but rather acts to reduce market uncertainty about what the standard is going to be, facilitating recombination. By providing standards for free, one also solves the λ problem and creates stronger incentives for further innovation. One may then see a role for nonprofit or government institutions in helping set standards.

A last comment regards possible market failures. Ex ante, if a bargaining problem (λ) stymies innovation, then one could integrate the firm and achieve the first best. With Bresnahan's starting point, however, the nature of knowledge distribution is such that one does not even know who to integrate with. That is the key problem: the fact that you cannot identify the recombinant possibilities ex ante means that you cannot easily solve the bargaining problem in practice—you cannot integrate your way around it. So innovation faces a serious market failure in the sense that socially profitable innovation does not occur. At the same time, it is not clear how a government realistically solves this problem directly, given that a government cannot obviously create a better information set (especially given the advantage of decentralized firms in perceiving innovative opportunities in their markets). Given the positive spillovers from the initial innovation, coupled with these fundamental information constraints, the government's role may then be limited to subsidizing innovation broadly—not just basic

---

[2] This consideration would suggest, for example, some value of the Bayh-Dole Act.

science, but also commercial innovation, through such policies as research and development tax credits.

In sum, this chapter points to knowledge distribution as a key feature in understanding innovation, with applications to general purpose technologies and other areas. This framework also points toward the tension between the openness that can allow recombination and the protection of one's own commercial interest that can incentivize the individual innovations themselves. In market settings, the profitability of the initial innovation will be sufficient for some innovative activity, and the market then acts to encourage recombination. In basic research settings, the institutions of public, open science can be understood in the same framework, but the analysis suggests that these science institutions may need a further look in helping to ensure that firms and publicly-supported researchers actually engage in efficient knowledge interchange.