The Impact of "Sticky" Information on Innovation and Problem-Solving

Eric von Hippel

WP# 3147-90-BPS April 1990

Support by the Sloan Foundation for the research reported on in this paper is gratefully acknowledged.

ABSTRACT

Researchers often regard information as being extremely slippery: perhaps expensive to generate, but with a marginal cost of diffusion (replication and transmission) near zero. Much of the special character of markets for information, and much of the difficulty in appropriating benefit from invention and innovation have been attributed precisely to this characteristic.

In this paper I build upon an alternate view that much information in fact has "sticky" qualities that make it hard or impossible to replicate and diffuse to many sites. I propose that sticky information can have a significant influence on patterns of problem-solving and innovation, requiring that problem-solving activity must shift to and among sites of sticky information as problem-solving proceeds. I offer examples and some empirical data that appear to support this hypothesis.

The Impact of "Sticky" Information on Innovation and Problem-Solving[1]

1.0: Introduction

Nelson (1980, 1982) has pointed out that economists have a rather schizophrenic view of information. On the one hand, they frequently view it as being extremely slippery stuff, perhaps expensive to generate, but with a marginal cost of diffusion (replication and transmission) near zero. On the other hand, at least since Polanyi (1958), there has also been a recognition that some information, such as "tacit" knowledge, may be very hard to diffuse.

The central tendency in economic theorizing has been to view information as slippery. Much ot the research on the special character of markets for information and the difficulty of appropriating benefit from invention and innovation has been based precisely on this characteristic. However, Nelson (1980, 1982), Teece (1981), and Kogut and Zander (1989), among others, have suggested that much information (data, knowledge) of interest to problem-solvers and innovators may in fact be costly to replicate and diffuse. There are a number of potential causes for this effect, and there are also important consequences that flow from it. At the moment, neither side of the issue is very well mapped or understood.

In this paper I assume that the (at least several) causes of information being costly to replicate and diffuse are independent of the consequences - at least,

[1]I would like to thank my close intellectual colleagues Anne Carter, Bradley Feld, Dietmar Harhoff, Rebecca Henderson, Wendy Mackay, Stephan Schrader, Sarah Slaughter, and Marcie Tyre for much discussion and advice on the ideas embodied in this paper.

"usefully often." That is, I assume that why one cannot have a unit of information will not influence the consequences of not having it. This assumption offers us the major convenience of being able to treat causes and consequences independently.

In order to clearly separate cause from consequence, I also propose that it will be useful to assign a term to represent the cost of replicating and transferring information independent of the reasons for such costs. In this paper I use the "stickiness" of information to refer to the cost of replicating and transferring information, where stickiness is a variable. At one end of the scale is information with zero stickiness, information that can be replicated and diffused at zero cost. At the other end is information that is so sticky that it cannot be replicated or diffused at any price.

In the paper I explore the idea that the stickiness of information can have an impact on the locus of problem-solving and innovation. In essence, I hypothesize that if information needed by a problem-solver is sticky, related problem-solving activity must move to and among loci of sticky information - perhaps repeatedly as problem-solving proceeds. The shiftings in the locus of problem-solving that I hypothesize can also affect the locus of innovation, since problem-solving is a major component of the innovation process. I provide examples and an existence test for the effect I propose, and I then argue that the effect will be found to be widespread. Not only are there many reasons for a given unit of information to be sticky but often, from the point of view of a problem-solver, the total costs of moving even minimally sticky information can be formidable - simply because a great deal of information is drawn upon to solve a problem. (For example, it may be economical to tell a home shopper who is engaged in problem-solving about any item in a store via telephone. It may even be economical to send out a particular item for home trial. But what if the shopper wants to see and handle "everything of type X in the entire store" before making a selection?).

I begin our exploration of sticky information with a brief review of a number of factors that can make information sticky (section 2). Next, I explore what I hypothesize to be an important consequence of sticky information: the shifting of problem-solving activity to and among sticky information sites. I examine two examples and also offer pilot data on the matter (section 3). I then review empirical evidence bearing on the possible ubiquity of sticky information (section 4) and discuss how the "stickiness" of information can be modified (section 5). Next, I propose that it is an economy, other things being equal, to design problems so that problem-solving draws on only one locus of sticky information, and I discuss how this is and might be done (section 6). Finally, I attempt to place the utility of the concept of sticky information in context and offer some suggestions for further research (section 7).

2.0: Sticky Information

Information may be sticky for a number of reasons. Some of these have to do with the volition of the information possessor, who may simply be unwilling to allow "proprietary" information to be replicated or transferred for reasons of private advantage. Often trade secrets, such as the source code for a computer program or the formula for a specialty chemical, fall into this category. Other factors have to do with how the information is encoded or embodied in a person, organization or object, with how it is indexed, and with the compatibility between information transmitter and receiver.

With respect to the difficulties for information transfer raised by encoding, consider that human skills are often "tacit" - not encoded in a way that lends to easy replication or transfer. Polanyi notes "the well-known fact *that the aim of a skillful performance is achieved by the observance of a set of rules which are not known as such to the person following them*" (Polanyi 1958, 49, italics in original). He illustrates by noting that swimmers are probably not aware of the rules they use to keep afloat (such as increasing buoyancy by not completely emptying their lungs when breathing out) nor are bicycle riders generally aware of the rules they follow in order to successfully stay on their bikes. Later, he argues that "an art which cannot be specified in detail cannot be transmitted by prescription, since no prescription for it exists. It can be passed on only by example from master to apprentice" (Polanyi 1958, 53).

With respect to the importance of indexing, one need only think about the vital role played by a card catalog in a library or about the utility of ordering the entries of a dictionary alphabetically. If an information base is not ordered in some way and if the indexing rules are not known to the information seeker, a perfectly encoded information set may be unusable.

With respect to compatibility between information transmitter and receiver, consider that information seekers must hold much related information if they are to hope to usefully "understand or operate" even unambiguous and well-indexed information that may be transmitted to them (Arrow 1974, 39). Thus, artists seeking to generate computer art using fractals will not typically be aided by the transmission of a math program containing that information. The recipients must either get the information they seek in "user friendly" form (which in practice means that the transmitter must understand what the recipients already know or can easily learn and must adapt access to the new information accordingly) and/or the recipients must learn the additional information needed to obtain and operate the information base.

Information can be embodied in natural or artificial products, sites, people, or organizations as well as in specialized information media such as magnetic tape. Sometimes one cannot or does not want to separate the information from its embodiment. In such a case the information can be either sticky or non-sticky, depending on whether its embodiment can be cheaply replicated and diffused to "compatible information receivers." For example, suppose that one cannot replicate and diffuse the information related to "how to run an automobile assembly plant efficiently" without replicating the embodiment of that information - the plant operating team and perhaps the plant itself. Such a replication would be difficult, and so that information would be "sticky" in our terms. On the other hand, one can easily replicate and diffuse the information embodied in a catalyst or computer chip by making and selling copies of those products. Therefore that information, although embodied, would be non-sticky in our terms.

3.0: Sticky Information and Its Effect on the Locus of Problem-Solving

Problem-solving involves operations applied to information. If the information a problem-solver calls for is not locally available, the traditional assumption of empirical researchers has been that the problem-solving individual or organization "acquires" it, that is, seeks it out and somehow transports it to where the problem-solving work is being done. Thus, Allen (Allen 1966; Frischmuth and Allen 1969) and Witte (1972) have empirically explored how problem-solvers acquire remotely located information.

But what if the required information is "sticky" and cannot cost effectively be replicated and/or moved to the problem-solver's site? In that case, I hypothesize that the locus of problem-solving itself will shift to the location(s) of that sticky information. And when multiple loci of sticky information are required to solve a given problem, I hypothesize that problem-solving activity will shift among such loci, perhaps repeatedly, as that sticky information is drawn upon during the course of problem-solving work. (In such instances, the outputs of problem-solving activities will be transferable, even though the information bases drawn upon are sticky. For example, an artist may not be able to transfer the entire information base that brings him or her to specify to a supplier, "I need a green paint of precisely X hue and luminance." However, it is quite simple to transmit the problem-solving result. Similarly, the responding paint manufacturer may be able to create and transfer to the artist the requested shade of green, but not be able to transfer the knowledge drawn on by that firm's chemists to achieve the feat.)

I first illustrate the hypothesis that problem-solving will shift to a single locus of sticky data via two examples. Then, I draw on a recent empirical study by Feld (1990) to provide a pilot test of the hypothesis that problem-solving activity will shift among loci of sticky data as problem-solving proceeds.

3.1: Shifting Problem-Solving to a Locus of Sticky Data: Two Examples

Let me illustrate the possible impact of sticky data on the locus of problem-solving (and innovation) via two examples. The first deals with a class of problem that lies within the personal experience of most readers - the planning of a business trip. The second considers patterns of problem-solving in two categories of semiconductor chips, DRAMs and ASICs.

Information drawn upon by a traveler to solve business trip-planning problems may come from many sources. For convenience, however, we will lump all possibly relevant information sources under only two headings here. One locus, the "traveler," actually comprises the traveler or a close proxy like an executive assistant, and those he or she seeks to visit and those people and events that have an influence on an individual's desire to travel. The second, the "supplier," actually comprises all firms and individuals who offer the services that may be used by a traveler in the course of a business trip, such as air transportation, hotels, and rental cars.

Information on the need for travel, preferred destinations, times, and constraints then are generated by and initially reside with the "traveler," while information regarding elements commonly used in the solution of travel planning problems is generated by and initially resides with trip element "suppliers," such as airlines. The problem-solving involved in the planning of a business trip requires access to elements from both of these information bases.

It is not difficult to observe that in today's world business trip planning is typically carried out at the traveler locus, with the traveler calling supplier-generated information to his or her locus as needed. Indeed, the latest trend is to give travelers direct computer access to online reservation systems containing supplier information on flights, hotels, rental cars, etc., thus eliding even a nominal problem-solving role by the supplier or travel agent (Gutis 1989).

Why does the traveler undertake the work of problem-solving in the instance of business trip planning? After all, on the basis of economies of specialization, one might think that trip suppliers have the advantage in this regard - major airlines or their agents deal with thousands of travelers daily who are struggling with this particular class of problem. Presumably, therefore, they have plenty of opportunity to generate specialized problem-solving expertise on the matter.

I propose that any economies of specialization that might attend problem-solving by suppliers are in this instance overwhelmed by the - I propose - sticky nature of relevant traveler information. Each traveler engaged in planning has a complex set of requirements for a satisfactory solution that is a function of a personal and complex information base. This information base is generated at the traveler's locus, changes rapidly ("If I change my trip to Tuesday, I can make an emergency visit to site X"), and is very rich in solution possibilities ("Uncle George lives in Phoenix; I wonder whether I can route through there on this trip?").

I suggest that the information just described meets our criterion for sticky information: it would be very difficult or perhaps impossible to encode, replicate, and shift to a remote location. In addition, it is so complex and poorly indexed for remote access that it appears impractical for a supplier to provide a traveler with a satisfactory solution via a series of precise inquiries directed to that traveler's sticky information in situ. (One can easily imagine the difficulties one might encounter in attempting such an exercise: "Do you want to travel? How about Phoenix? Does my query remind you that you need to consider visiting a supplier in Detroit? Well then, what about Detroit? On second thought, do you necessarily need to travel there to solve the problem that I just evoked?")

In sum I suggest that, in the instance of business trip planning, the traveler's information base is "sticky" while the supplier information base is not. In turn this results, I propose, in the siting of this category of problem-solving activity at the traveler locus.

If the above reasoning as to the importance of the locus of sticky information to the locus of problem-solving is correct, we should be able to propose conditions under which relevant traveler information is not sticky and see the locus of travel planning change - and we can. Sometimes a traveler may have a very simple, non-contingent set of personal requirements relevant to a particular trip. For example, he or she may want "a two-week tour of Europe during August." Under these circumstances, it would appear efficient for the traveler to transfer this relatively simple need information to the supplier and benefit from the supplier's economies of scale with respect to detailed planning of destinations, transport and hotel accommodation, sightseeing tours, etc. And, "vacation travel packages" appear to conform to this expectation. Here, suppliers rather than travelers are the locus of detailed travel problem-solving.

Let me now quickly add a second case to enrich our discussion. The content of this case will not be so intuitively clear to most readers, but the case does offer the advantage of showing that the issues discussed in the instance of travel planning apply equally to cases of "real innovation."

My second illustration deals with Dynamic Random Access Memory semiconductor chips, commonly referred to as DRAMs. DRAMs are used today as the memory devices in electronic devices ranging from hand-held calculators to mainframe computers. Need information regarding the function that these chips should perform are generated by end users of systems that contain DRAMs and by the designers of such systems - whom we will also call "users" for the moment. However, information bearing on how one might design and produce chips that are responsive to user needs is generated at the locus of chip manufacturers and their suppliers - "manufacturers."

In contrast to the pattern in the planning of business trips, the pattern in the instance of DRAM design is to transfer need information to the manufacturer and to conduct innovation-related problem-solving at the manufacturer locus. I propose that this information transfer pattern, too, appears reasonable in the sense that manufacturer capability information relevant to DRAM design appears sticky, while user need information does not. Consider that the design of DRAMs and related fabrication processes represents the very front edge of the chipmaker's technology and art. The manufacturer's information base relevant to these problems is very rich and complex, involving many design and process elements and trade-offs. This information base changes rapidly ("On the basis of yesterday's experiment, I now understand that I can do X") and, at any given time, key portions are so new and dimly understood as to exist only in the form of intuitions in the minds of creative engineers. In sharp contrast, the need information generated in the user locus relevant to DRAM design is relatively simple and can be encoded in a few unambiguous parameters such as pulse shape, speed, cost, size, energy consumption, etc.

As in the earlier travel-planning example, if our reasoning as to the importance of the locus of sticky information to the locus of problem-solving is correct, we may see the locus of problem-solving shift from user to manufacturer in the instance of semiconductor chip design - if we can find conditions where user rather than manufacturer information appears to be sticky. I propose that the conditions surrounding the design of specific ASICs (Application Specific Integrated Circuits) meet our criteria.

An ASIC is an integrated circuit that is customized to the needs of a particular user. Specification of its function requires an understanding of a rich and richly constrained information base originating at the user locus having to do with the system of which the ASIC will be part. In contrast, the supplier capability information required by a user who wishes to design an ASIC involves a set of standard rules and constraints that guides the user and prevents him from doing designs that exceed the capability of the manufacturer's silicon foundry process. Today, supplier-developed ASIC design rules have been made easily transportable to the user locus: ASIC manufacturers have worked to embody all the user needs from the supplier information base in "user friendly" CAD packages that the user can operate at the user site.

Thus, today, we see that the user information base relevant to ASIC circuit design appears to have attributes of stickiness, while the supplier information base relevant to the design of specific ASICs does not. And, as we would predict, we also see the design of custom ASIC circuits being carried out at the user locus rather than at the locus of the integrated circuit manufacturer. (Note that in this example ASIC suppliers invested in <u>converting</u> elements of their information base into a form that could be transported to user sites. We will develop this point in section 5.)

3.2 The Shifting of Problem-Solving Activity
Among Loci of Sticky Data: A Pilot Test

My second hypothesis was that when multiple loci of sticky information are required to solve a given problem, problem-solving activity will shift among such loci, perhaps repeatedly, as that sticky information is drawn upon during the course of problem-solving work. Let me next draw on some very interesting empirical findings by Feld (1990) to provide a pilot test.

Feld examined the problem-solving interactions between a firm in the business of developing applications software and three of its clients during the course of three unrelated development projects. Interestingly, he found that "rich" information tended to stayed fixed at its site of origin during the problem-solving process, while problem-solving activity repeatedly shifted back and forth between user and manufacturer sites. Let me briefly outline how this finding was obtained.

At the start of each of the three software development projects studied by Feld, the software firm and the single client engaged in that project with the firm each had rich information bases that were needed for problem-solving and that were unknown to the other. In each instance the software development firm had rich software development expertise unknown to the client. And in each instance the client had rich need information unknown to the development firm. (In order to insure that this was so, each of the three clients examined was in a different business - specifically, venture capital, insurance, and group medical practice - and also was the first clients in their business to be served by the programming house: At the start of each programming project examined, it was clear that the software house had no expertise in the business fields of these clients.)

It was the software development firm's practice to progressively extend and improve the software produced for a client by means of periodic revisions. Each revision incorporated a number of independent changes, and the firm kept contemporaneous records of each revision and of the nature of the changes incorporated in each. Feld used these records plus interviews with client and software firm personnel to determine the nature of the information passed between client and software firm during each of 516 changes incorporated in a total of 68 revisions made to the three projects studied.

For our purposes, Feld's key finding was that rich information on need and capability were not transferred between software firm and clients in the instances studied - even after several years of repeated interaction. Instead of finding rich mutual interaction between client and developer behind the revisions and requests coded, he found very brief interactions instead.

Most (81%) of the 516 changes in Feld's sample were requested/suggested by the client rather than the developer, with the rest being suggestions from the developer to the client. In the instance of client-to-developer requests, the content of the request was, in effect, a simple "Do this," with no accompanying explanation or request for discussion. The response of the software firm to such requests was equally spare and unadorned: Simple notes were attached to revised software sent to the client saying, in effect, "We did what you asked." No information was provided (or, apparently, sought) regarding the context or any explanation of the reasoning used to develop the particular response. In the instance of developer-to-client suggestions, the content was equally spare and consisted of simple "If you do this, it will be more convenient for you" suggestions. Importantly, these developer suggestions were derived from information specific to the developer's knowledge base rather than from client-specific knowledge. For example, the developer might suggest to a client that "[We know from our understanding of software that] The software will run faster if you do X."

In Table 1, we see this finding elaborated. Feld divided his sample of specifications for changes into two categories, "chunks" and "items." Both were very specific requests, but chunks were somewhat more complex and required somewhat more programming than did items. For example, a client request for a chunk would be a request to "Generate a table that looks like this." In contrast, a typical request for an item would be: "Where you now print 'share' in the output reports, print 'shares' instead."

Feld also coded the content of the requests into two categories, "complete" and "incomplete." In the former case, the initial "Do this" request was complete and unambiguous enough to allow implementation without further information; in the latter case the "Do this" request needed further amplification before it could be implemented.

Content of Messages
Between Developer and Client

	Number of "Do This" Messages			
	Complete		Incomplete	
Project Studied	Chunk	Item	Chunk	Item
Venture Capital	6	78	16	2
Insurance	10	118	17	2
Medical	25	211	28	3
TOTAL	41	407	61	7

Table 1: Messages Passing Between Software Developer and Client Regarding Software Changes. (Messages passing between software developer and client with respect to software changes were found to consist of brief "Do this" messages that asked for or suggested changes that varied in terms of the amount of programming required for implementation [coded as "Chunk" vs. "Item"]. The messages also varied in terms of their completeness [coded as "Complete" vs. "Incomplete"]. [Feld 1990, Table 5.])

As can be seen from Table 1, "Do This - Incomplete" requests were almost entirely found in the instance of chunks. In such instances the developer would seek further information from the user. (All "Do This - Incomplete" requests were from client to developer.) Requests by the developer for more information were found to be of a "How do we do what you want to do" nature, rather than a query as to why the client might want to do what was requested. For example, the developer found the following request from a client to be incomplete: "Use a percentage of cost method for revaluation calculations." The additional information the project manager sought in this instance: "How does one do the calculation you request?" No information was sought from the client that would enable the software company to understand the purpose or merit of the client request. Nor did the firm attempt to generate such information from other sources. For example, in the instance of the request just noted, the software company did not take it upon itself to learn about cost valuation equations in order to understand the client request better.

In the case of the venture capital project, the client did all the need information analysis that led to the generation of these requests. In the case of the insurance and medical group practice projects, software firm personnel sometimes traveled to the user site to gain access to user information needed to frame the "do X" statements to be transmitted back to the software house. Rich user information was not carried back to the software firms as a result of these visits, however. Feld found that software firm personnel who made the visits reported that they concentrated narrowly on "figuring out what we were supposed to do" rather than on acquiring a broad understanding of the clients' business or even an understanding of the context of the particular need being framed.

The generality of the pattern found by Feld in software is suggested by the pattern of trial and error that is often encountered during product and service development. This can involve repeated alternating calls on need and solution information by problem-solvers. If the need and solution information bases are tied to different sites, we should see a pattern in which problem-solving activity shuttles back and forth between these two sites. First, as shown schematically in Figure 1, a problem-solver may draw on a user need information base to generate some attributes for a desired new product or service.

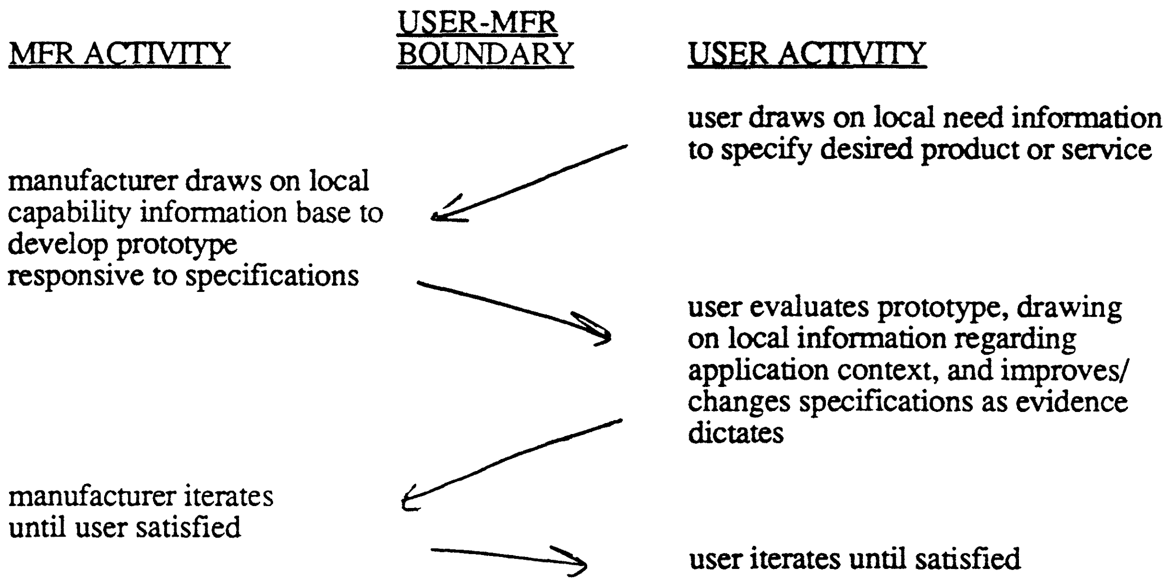

Figure 1: Trial-and-Error Pattern Often Encountered in Product and Service Development. (In this pattern, problem-solving activity shuttles between a user need information base and a supplier capability information base.)

Then, a manufacturer information base may drawn upon in order to develop a prototype that appears responsive to the specification. The prototype is then tested ("trial") within its proposed use context to verify function and the accuracy of the initially stated need. If the two do not match satisfactorily - and they often do not ("error") - need and/or capability information must be revisited in search of a closer match. This cycle may be repeated few or many times until an acceptable match is found.

I propose that a pattern in which problem-solving shuttles repeatedly between user and supplier information base locations is an indicator that all information relevant to solving the problem at hand have not been transferred from one or both sides to the other. Full transfer of all relevant need information from user to manufacturer should allow a manufacturer to make a responsive product and test its appropriateness within the manufacturer plant - no additional problem-solving at the user site would be necessary. Similarly, full transfer of all relevant capability information to the user should allow the user to solve his need within his own firm without a need for repeated interaction with the supplier.

Interestingly in this regard, in recent years software developers have developed an explicit pattern of problem-solving activity that in fact does shuttle repeatedly between user and manufacturer loci. Many have now implemented the pattern under the name "rapid prototyping." Prior to the introduction of this procedure, it was standard practice for software developers to meet users and agree on a product need specification at the start of a project, and then to work isolated from further user contact until the product was delivered. However, via painful and costly experience, developers found that users often rejected the products developed via this method. Users said, in effect, "Although this software does meet the initial specification we both agreed upon, I now find upon experimentation that it is <u>not</u> what I need."

In eventual response, software engineers developed rapid prototyping methods[2] that allowed the simulation of key functions of proposed software products. Users could then experiment with prototypes representing what they initially thought they needed, so that they could refine their need specifications in

---

[2]The key technical advances in programming techniques that made rapid prototyping possible were the development of application generators oriented around data-base management systems, plus the development of report-generation systems (Davis 1986).

the light of what they learned. This process of alternate problem-solving in the locus of need information and the locus of capability could then proceed as often as needed to get a good fit between need and solution. Today, the rapid prototyping approach has been tested by many expert software designers and typically has been reported to be far superior to the old approach with respect to accurate understanding of user need (Hekmatpour 1987; Zelkowitz 1980; Klausner and Konchan 1982; Gomaa 1983). The single experimental comparison of the two methods I am aware of (Boehm, Gray, and Seewaldt 1984) also comes to that same conclusion.

In sum, we have observed patterns of problem-solving in which problem-solving activity is shifted to the locus of rich information as that information is needed. Does this mean the information at issue is in fact sticky? After all, the pattern may simply represent a choice by the problem-solver to shift the locus of problem-solving instead of shifting needed information under circumstances where both possibilities were available.

Recall that information can be tied to a particular site for proprietary and/or technical reasons. In the instance of the Feld study, we found that the manufacturer saw a proprietary advantage in keeping its software development capability at the manufacturer locus as a trade secret. Thus, problem-solvers located at this firm's clients clearly did not have a choice in this instance: They were forced to shift problem-solving to the software development firm if they wished that firm's sticky information to be applied to their problem. The matter is less clear in the instance of client information. Feld reported (personal communication to the author) that neither side considered the possibility of transferring very rich client information to the software development firm. Therefore, issues of proprietary advantage and technical feasibility that would affect the stickiness of client information were not explored in this instance.

In the experience that led software developers to develop rapid prototyping, however, I think that we can infer the existence of sticky client information. The literature in that field, typically taking the viewpoint of the software developer, does not refer to client unwillingness to transfer need information to inquiring software developers. Instead, it refers to the inability to successfully transfer all useful information regarding user need to the would-be supplier at the start of a project, often despite clear awareness on both sides that successful transfer of this information was critical to overall project success.

4.0: The Possible Ubiquity of "Very Sticky" Information

At the start of this paper, I argued that information would not have to be very sticky to have an impact on the locus of problem-solving - if access to a lot of information (and/or to rapidly changing information) was required to solve a problem. Nonetheless, it is interesting to note some evidence from natural experiments in a number of fields that collectively suggest that "very sticky" information may in fact be ubiquitous. Let me therefore consider some findings in the fields of expert systems and marketing research from this point of view. In both of these fields, a major goal is precisely the replication of information located at one site and the shifting of this information to other sites. And, in both fields, this goal has been found to be very difficult to achieve.

4.1: Sticky Information and Expert Systems

Expert systems are a type of computer program designed to "capture" and replicate the expert problem-solving abilities of human specialists in a given field. Developers of expert systems have a strong incentive to learn to encode and diffuse the information held by such specialists, because their research has clearly shown the central role that information plays in most human expertise. (Indeed, many now speculate that the essence of problem-solving expertise lies as much with an expert's possession of rich domain-specific information as with his or her possession of especially sophisticated problem-solving skills [Davis 1986; Simon et al. 1987; Larkin et al. 1980; Feigenbaum 1978].)

Many expert systems have been successfully developed and applied. Thus, a number of medical expert systems have been developed to aid doctors in the diagnosis of some types of disease. (For example, an expert system called PUFF has been found a useful assistant in the diagnosis of some lung diseases.) In other fields, Digital Equipment developed a very successful system that aids in the proper specification of interconnections between computer system components; American Express has developed a very effective system to aid in the evaluation of credit risks; and so forth.

Recently, improved software tools have made it much easier for users to develop expert systems without extensive aid from specialists, and many users are responding to this opportunity. A typical example is an expert system developed by a product-testing group at IBM to capture the knowledge of technicians within the group who were especially expert in diagnosing problems in a specific model of computer disk drives (Feigenbaum, McCorduck, and Nii 1988).

Although they have had many successes, specialists who develop expert systems have learned that, today, there are some kinds of human expertise that they <u>cannot</u> encode, replicate, and transfer successfully. I propose that this observation is useful evidence for the stickiness of a major class of real-world information. Let me first develop the basic observation a little further, and then return to the matter of stickiness.

The expert knowledge component of an expert system is embodied in the form of a series of rules. (Thus, the expert knowledge contained in the IBM diagnostic system mentioned just above consists of about 2,000 "IF - THEN" rules [e.g., "IF you see X symptom, THEN look for Y type of fault"] [Feigenbaum, McCorduck, and Nii 1988, 69].) It has been found by those who build expert systems that experts typically do not themselves encode their knowledge in the form of explicit rules for their own use. Therefore, their knowledge must be converted into such rules before it can be installed in the system. The conversion of expertise into explicit rules is not an easy or certain task. In essence, experts in a subject matter, with or without the aid of specialist "knowledge engineers," seek via trial-and-error to devise rules capable of producing outcomes similar to those they produce via their own expertise.

Today, only a very fragile and limited simulacrum of human expertise can be embodied in an expert system. (One of the fragilities noted by Randall Davis[2]: An expert system is typically unable to determine whether a posed problem is appropriate to its particular domain of expertise. For example, as Davis amusingly points out, if you tell a medical system whose domain of expertise is antibiotics that you are having trouble with your bicycle, it will try its best to solve the problem with an antibiotic.)

Importantly for our purposes, knowledge engineers have learned by experience to judge when efforts to encode knowledge into a useful expert system are likely to be successful. First, they find, the human expertise to be encoded in a proposed expert system should involve a sharply delimited area (e.g., "trouble shoot X model disk drive" rather than "trouble shoot all disk drives"). Second, the expertise should not require one to draw on "common sense." Third, the expertise to be encoded should be "one you can talk about over the telephone" rather than, say, a physical skill like swimming.

The first two rules of thumb deal with the recognition that many types of human expertise cannot be simulated via the methods used in today's expert systems simply because the number of rules required is too large. (For example, the things humans know about the world that expert developers term "common

---

2 In my discussion of expert systems I gratefully acknowledge the insights provided to me by Randall Davis, Associate Director, Artificial Intelligence Laboratory, Massachusetts Institute of Technology. Also see Gevarter (1985, 55-58).

sense" may encompass millions of complexly interrelated items. (E.g., "When it rains, I may get wet if I am outside - unless I have an umbrella and it is not windy, or unless I have a raincoat if it is windy, or unless I am with someone who has an umbrella, or unless I am on the lee side of a building, or unless....") The third rule of thumb simply says that there are whole classes of human expertise that one does not yet know how to encode in the form of rules.

Taken together, I propose, the three experience-based rules just mentioned indicate that much human expertise is effectively sticky today: Such expertise can be replicated and diffused to remote sites via an apprenticeship process in which experts train novices, and then diffuse the information by physically transferring the new experts. But this process can be long and costly and of uncertain effectiveness. (Note that simply transferring the original expert to a remote site does not involve the replication of sticky information.) In contrast, as noted by a member of the team that developed the disk drive diagnostic expert system: "With the expert system, we know that once we have the knowledge captured, it's easy to transfer. We've demonstrated that we can do that very well. San Jose is the primary manufacturing center for this device, but we have sister facilities in Japan, Germany and Brazil. We shipped our knowledge bases to them and we've instilled our methodologies in their organizations" (Feigenbaum et al. 1988, 68).

4.2: Sticky Information and Marketing Research

A second area where great efforts have been made to encode and transfer information felt to be very valuable - with, however, only limited success - is in the field of marketing research. I suggest that, here as in expert systems, the combination of a strong effort to encode and transfer information in conjunction with limited success is an indicator that the information involved is "sticky." Let me fill in the case a little further so that the reader will have some independent basis for assessing this claim.

I begin by noting that one third to one half of new products introduced to market are commercial failures (see studies cited in Urban and Hauser 1980, 3) and that the most important cause of such failure is a poor understanding of user needs (Rothwell et al. 1974). This situation persists despite many years of effort by academic and industrial specialists in marketing research to develop better ways to apprehend and encode user-based need information and to transfer such information to manufacturers' sites. Why is this problem so difficult?

A manufacturer needs a great deal of information to accurately understand user need for a product or service that differs substantially from existing ones. This is because individual products and services are only components in larger usage patterns that may involve many such. Since a change in one component can change perceptions of and needs for some or all other products in that pattern, an accurate understanding of the need for a new product or service must include an understanding of existing multiproduct usage patterns in which the new product might play a role. Then, one must be able to collect the data that would allow one to understand and evaluate the new product's potential contribution to these and/or to new patterns that the proposed new product may make possible and attractive. (For example, a change in the operating characteristics of a computer may allow a user to solve new problem types if he makes related changes in software and perhaps in other, related products and practices. Similarly, a consumer's switch to microwave cooking may well induce related changes in food recipes, kitchen practices, and kitchen utensils.) Finally, since substitutes exist for many multiproduct usage patterns (e.g., many forms of problem analysis are available in addition to the novel ones made possible by a new computer), the manufacturer must gather the user data relating to how the new possibilities presented by the proposed new product will compete (or fail to compete) with existing options.

The information collection task just outlined is clearly a daunting one, and in practice it is not done. Instead, marketing research methods have followed the route of attempting to make do with only a very little information from users. Consider, for example, the interesting class of "multiattribute" marketing research tools. These are designed to analyze user perceptions of products and their preferences regarding them. The information actually collected from users to achieve this purpose is quite limited, however.

Essentially, users are asked to identify approximately twenty scalable attributes that they associate with products in a category under study. (For example, users of cameras may report to researchers that "picture sharpness" and "camera weight" are two important "attributes" of cameras from their point of view.) Then, users are asked to indicate the importance of each attribute to them and to indicate their preferences for a range of existing products in the category. This information is then used to analyze user preferences for proposed new products. (This is done by analyzing products as bundles of attributes that have been specified and weighted as to desirability by users.)

Such methods represent the current apex of the market researcher's art. But researchers in the field are clear about the method's limits. Thus, "the . . . framework is primarily useful for locating 'new' product opportunities which may not be substantially different from current alternatives" (Shocker and Srinivasan 1979). And, "the applicability of the system is limited to situations in which the new brand seeks to penetrate a product category well-defined in terms of the nature and closeness of substitutes. Cases in which a very novel or innovative offering effectively creates a new product category cannot be handled by these methods" (Silk and Urban 1978).

The reason for these limits is clear, I propose. As noted earlier, the methods actually transfer very little of the (sticky?) information that may be relevant to user needs for a new product or service from the user to the manufacturer.

5.0: The Modifiability of "Stickiness"

Earlier, we noted that information might be effectively sticky because information owners have sought to protect their private advantage (as in trade secrets), and/or because the information is poorly encoded for transfer purposes (as in human skills), and/or because the information is poorly indexed, and/or because incompatibilities exist between transmitter and receiver. Interestingly, none of these factors is immutable. All can be affected, at least to some degree, if expectations of benefit exist to bring forth the investment required to do so.

Considerations of when to invest in lessening the stickiness of information should be similar to those affecting any innovation investment. The economic costs of gaining access to sticky information consist of (1) property rights-related costs and (2) technical and production costs. Costs associated with property rights considerations would, for example, include payments to owners of proprietary information to induce them to release such information. Technical and production costs would, for example, in the case of probing a given plot of land for geological information, include the costs of developing and/or obtaining sensors and seismic measuring equipment and operating them to obtain information that could then be coded for remote storage and analysis. Similarly, one may invest in learning, or attempting to learn, to encode the skills of swimmers and bicycle riders, or the collective behaviors of organizations, for purposes of analysis or replication at a remote site. (The development of expert systems is based precisely on such work.)

Stickiness due to poor indexing can also be affected by investment in technological advances. For example, an unaided human would find it very difficult to use a dictionary having randomly ordered entries. However, a human equipped with a high-speed computer might find a randomly ordered dictionary as convenient as an alphabetically indexed one: Each time one wished to look up a word, the computer could simply be instructed to scan every dictionary entry until the right one was found. And finally, given the proper investment, it may be possible to endow a would-be transmitter of information and the would-be receiver with the additional information needed to allow the exchange of information and to enable its remote use.

In the case of travel planning, discussed earlier, we saw that the information that the business traveler needs from the supplier information base for planning the typical trip is today offered online via a convenient menu of clearly described options regarding trip schedules, car rentals, hotel locations, etc. We should now note that it took a great deal of investment by the travel industry to reduce the stickiness of this information. Thus, airlines have invested many millions to develop the online information bases that now allow passengers as well as ticket agents and travel agents to conveniently check flight availability and costs and to book passage on particular flights.

After weighing the costs involved, those deciding whether to invest in reducing information stickiness in a given situation would also consider how they might expect to appropriate returns, with many situation-specific matters entering into their calculations. For example: Does X specific type of information, once analyzed/processed/encoded for transfer, become a free good? Or can one hope to profit from one-to-a-customer information sales (as in information encoded in books)? Or can one hope to profit from repeat information sales to the same customer (as with the ever-changing information in airline information bases)? Or, can one hope to profit from the sale of goods or services associated with the information (as in the sale of airline flights)? And so on.

Note that such calculations might show a net gain from an investment in reducing stickiness - or a loss. For example, information possessors may find benefits obtainable by making information transferable may be outweighed by benefits obtainable via maintaining stickiness under some circumstances. Thus, a semiconductor user may gain in problem-solving efficiency by providing suppliers with access to his "need information base." At the same time, however, he may lose if would-be competitors gain early knowledge of his plans through this action.

6.0: The Advantage of Needing "A Maximum of One" Locus of Sticky Information For Problem-Solving

I propose that problem-solvers gain an advantage by concentrating problem-solving in a single locus, "other things being equal." My reasoning is that the passing back-and-forth of problem-solving activity as necessary to access multiple sticky information bases is an additional cost that is not present in instances when all problem-solving activity and information called on by problem-solvers are located at a single site. I find no quantitative information bearing on this matter, but common-sense observation suggests that the additional cost can be considerable. Consider, for example, the enthusiasm with which users endorse (and adopt) single-site "desk-top publishing" over earlier practice involving repeated back-and-forth transactions between user, graphic designer, and printer. Similarly, consider the economy and convenience of home medical diagnosis kits, for example, glucose detection kits, relative to earlier practice that involved coordinating transactions between information tied at three sites - a patient, a doctor, and a medical laboratory.

Suppose that for the moment we accept that there is in fact an economy to be gained by concentrating problem-solving in one site rather than distributing it among two or more, "other things being equal." Suppose also that the encoding of information bases for replication and diffusion represents an investment in addition to the investment that must be made to use a sticky information base at its present locus. Given these two assumptions, we come to the interesting hypothesis that it would be useful to strive to have sticky information relevant to the solution of a given problem at only one site if any - with that one site being the same as the preferred site of problem-solving activity.

Assuming that one can identify in advance the information that a problem-solver will call on during problem-solving (and, elsewhere, I argue that one can in fact often do this [von Hippel 1990]), there are two complementary approaches to achieving the end of having only one locus of sticky information relevant to a given problem. (Of course and alternatively, one can also change the problem and/or one's specifications for what constitutes an acceptable solution.) First and straightforwardly, one can invest in reducing the stickiness of relevant information residing at all sites except the preferred site for problem-solving activity. Second, one can attempt to "partition" the problem at issue into subproblems that each depend on only one locus of sticky information. Let me illustrate this second approach by reference to the examples of travel planning and semiconductor chip design presented in section 3.

In the travel planning example discussed earlier, the fact that the individual business traveler was the locus of problem-solving with respect to his or her individual trip did not imply that the supplier (say, the airline) was not involved in any related problem-solving. Clearly, it is important to have airline flights scheduled so as to be convenient for the business traveler. And, clearly, the matter of optimal scheduling also must draw heavily on sticky supplier information having to do with efficient aircraft routing in the light of maintenance needs, the characteristics of available airplanes, etc..

The travel problem could have been framed as one requiring sticky information residing at both user and supplier sites for solution. Instead, it has been partitioned into two subproblems - each soluble by reference to only one site of sticky information. Thus, when the traveler attempts to solve an individual trip-planning problem, the airline offers him only a fixed set of options and thereby decouples its problem-solving from his. (An airline will not add another flight to Phoenix if an individual passenger requests it.) And, when the airline does engage in problem-solving with respect to flight scheduling, it draws only on sticky information residing at its own site. (The only information it calls for from passengers has to do with encoded and transferable information on aggregate flight choices. It does not call on the sticky information behind each individual choice that resides at user loci.)

Similarly, the traditional means of providing custom integrated circuits (ASICs) to buyers involved a shuttling between two loci of sticky information in the manner shown schematically in Figure 1, earlier. The user would first draw on sticky need information to evolve specifications for the circuit needed and transfer this specification to the supplier. The ASIC manufacturer would then attempt to create an integrated circuit responsive to the specification by drawing on sticky information resident at his site bearing on semiconductor design and processing. Then, the manufacturer would transfer a prototype to the user for testing, and the shuttle would continue until an acceptable match between need and solution was achieved.

The genius - and popularity - of the ASIC form of custom circuit design and manufacture is the partitioning of the overall problem into two sub-problems. Each draws on only one locus of sticky information, thus eliminating the need to shuttle between two such sites. The manufacturer of ASICs draws on his own sticky information to develop and improve the function of his "silicon foundry" and also works on generic circuit precursors (e.g., a "sea of gates" design) independently of any user request for a specific circuit. He then provides the user with the general rules one can use to design a specific ASIC while staying within the capabilities of the processing plant - the "silicon foundry" - via a user-accessible CAD package. The user then draws on this non-sticky information from the manufacturer and on his own sticky information to develop the specific circuit he needs.

7.0: Discussion

Sometimes we have a trained inability to see certain phenomena because of the way we are used to asking questions. I propose that this is the case with respect to the likely ubiquity of sticky information.

First, note that researchers interested in the costs of information transfer have tended to focus on the implications of easy transfer. Thus, the focus of intellectual property law is essentially on how to preserve property rights in the instance of information that can be transferred. Similarly, economists interested in markets for information have tended to focus on the economic implications of information that may be costly to create, but that can be replicated and transferred at low cost. (Thus, Arrow points out, "the cost of transmitting a given body of information is frequently very low. . . . In the absence of special legal protection, the owner cannot, however, simply sell information on the open market. Any one purchaser can destroy the monopoly, since he can reproduce the information at little or no cost" [Arrow 1962, 614-15].) Such efforts are certainly not misplaced, but the underlying assumption with respect to the ease of transferring information may have had the inadvertent effect of diverting attention from the fact that much (most?) information may in fact be sticky.

Second, observe that researchers interested in patterns of information transfer such as Allen (Allen 1966; Frischmuth and Allen 1969) and Witte (1972) have focused their empirical research and interest on how problem-solvers "acquire" remotely located information by moving it to the site where they happen to reside. These researchers did not, however, explore how problem-solvers and/or problem-solving activity may or may not shift to the site of information called for by problem-solvers. The inadventent effect of the focus chosen may have been to divert attention from the stickiness information can sometimes possess.

Third, note that those interested in human problem-solving have, until quite recently, tended to focus on problem-solving "strategies" that might be employed by problem-solvers, rather than on the information that they might require. Indeed, empirical research on the topic was often designed to "control out" the information problem. This was typically done implicitly through a focus on problems for which all needed information was (assumed to be) already in the possession of experimental subjects or was provided to them as part of the experimental setting. Luchins's "water jar" problems are such an example. (These had the general form: "Given water jars with a capacity of A, B, and C quarts respectively, how can you measure out D quarts of water?" [Luchins 1942].) The crypt-arithmetic problems used by Newell and Simon (1972) also fit this mold. Again, the inadvertent effect of this emphasis was to divert interest and attention from the attributes of information used in problem-solving, including the attribute of stickiness.

Recently, however, this situation has begun to change, at least with respect to the way a significant number of psychologists interested in human problem-solving frame their research. As noted earlier, researchers in artificial intelligence are becoming very interested in the encoding of information required for the development of expert systems. In addition, some psychologists have begun to explore "practical intelligence" (Scribner 1986; Scribner 1984).[3] Here, the importance of the context of problem-solving - and related rich information - has been made clear in the instance of many every-day tasks such as grocery shopping, meal-planning, and errand running (Lave, Murtaugh, and de la Rocha 1984; Hayes-Roth and Hayes-Roth 1979; Byrne 1977). We can expect issues of information stickiness to fit well with the interests of researchers in these fields.

---

3 My thanks to Katherine Nelson, Professor, City University of New York, for introducing me to this very interesting field.

7.1: Suggestions for Further Research

Considerations of information stickiness can have an impact on a number of important areas of research that deal implicitly or explicitly with the cost of replicating and/or transferring information. It would be useful, I think, to explore and develop the implications of sticky information for these. For example, consider that patterns in diffusing, selling, protecting, trading, and appropriating benefit from information should be strongly affected by the presence or absence of stickiness. Thus, innovations dependent on sticky information for diffusion would presumably show different diffusion patterns than innovations not having this characteristic. For example, they might not diffuse at all, or they might diffuse in a pattern dictated and enabled by the information holder. (Suggestively in this regard, Rogers [1983, 231] cites three studies of farm innovations that find that an innovation's "complexity" [perception by the user of the relative difficulty of understanding and using an innovation] slows adoption.)

In addition, further explorations of the real-world causes of information stickiness would, I think, be most valuable. I have listed a few such causes in section 2 (encodability, etc.), but we need empirical evidence on the relative importance of these in the real world. Also, the list of causes may be incomplete in important respects - we just do not know at present. One might begin research into these matters with hypotheses about the characteristics of information that will make it easy or hard to transfer, such as the two dimensions of codability and complexity proposed by Kogut and Zander (1989, 6-10). And/or one might simply begin with "grounded research" (my own preference) into the causes of stickiness in some few categories of innovation.

Another interesting research area, I propose, would explore the implications of sticky information for the innovation process. At the gross level, if sticky information is involved in solving an innovation-related problem and can affect the cost of problem-solving, it is reasonable to think that there may sometimes be an impact on "the" locus of innovation (von Hippel 1988). At a finer-grained level, the locus of sticky information might help to explain patterns of joint development such as those observed between users and manufacturers of products in some fields. (Thus, Shaw [1986] and Biemans [1989] have shown the extensive interaction between medical specialists and equipment makers in the instance of medical equipment development.) Similarly, one might look at the possible impacts of sticky information on patterns of cooperation observable between manufacturers and their suppliers, such as those explored by Clark and colleagues in the United States, European, and Japanese auto industries (Clark, Chew, and Fujimoto 1987).

Further, considerations of sticky information may also help to explain when and why factory workers will have advantages over more expert engineers with respect to problem-solving and innovation regarding at least some aspects of processes they use (Zuboff 1988). (Here, I would propose that workers have the advantage when they are in closer contact with the information needed for problem-solving, and when that information is very sticky - or somewhat sticky but rich and/or rapidly changing.) And, considerations of sticky information, which characterize the information a firm can gain access to, may serve as a useful complement to research that explores whether and how firms may innovate as a function of how they "filter" information (e.g., Henderson and Clark 1990).

References

Allen, Thomas J. "Studies of the Problem-Solving Process in Engineering Design." IEEE Transactions on Engineering Management EM-13, no.2 (June 1966):72-83.

Arrow, Kenneth J. "Economic Welfare and the Allocation of Resources of Invention." In The Rate and Direction of Inventive Activity: Economic and Social Factors, A Report of the National Bureau of Economic Research, 609-25. Princeton, N.J.: Princeton University Press, 1962.

Arrow, Kenneth J. The Limits of Organization. New York: W.W. Norton, 1974.

Boehm, Barry W., Terence E. Gray, and Thomas Seewaldt. "Prototyping Versus Specifying: A Multiproject Experiment." IEEE Transactions on Software Engineering SE-10, no. 3 (May 1984): 290-303.

Byrne, Richard. "Planning Meals: Problem-Solving on a Real information-Base." Cognition 5, no.4(December 1977):287-332.

Clark, Kim, W. Bruce Chew, and Takahiro Fujimoto. "Product Development in the World Auto Industry: Strategy, Organization and Performance." Paper presented at the Brookings Microeconomics Conference 3-4 December 1987.

Davis, Randall. "Knowledge-Based Systems." Science 231, no.4741(28 February 1986):957-63.

Feigenbaum, Edward A., "The Art of Artificial Intelligence: Themes and Case Studies of Knowledge Engineering." Paper presented at the 1978 National Computer Conference, 5-8 June 1978, Anaheim, Calif. In AFIPS Conference Proceedings: 1978 National Computer Conference, edited by Sakti P. Ghosh and Leonard Y. Liu, 227-40. Montvale, N.J.: AFIPS Press, 1978.

Feigenbaum, Edward, Pamela McCorduck, and H. Penny Nii. The Rise of the Expert Company: How Visionary Companies Are Using Artificial Intelligence to Achieve Higher Productivity and Profits. New York: Times Books, 1988.

Feld, Bradley A. "The Changing Role of the User in the Development of Application Software." Working Paper (forthcoming). Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Mass., 1990.

Frischmuth, Daniel S., and Thomas J. Allen. "A Model for the Description and Evaluation of Technical Problem Solving." IEEE Transactions on Engineering Management EM-16, no.2 (May 1969): 58-64.

Gevarter, William B. Intelligent Machines: An Introductory Perspective of Artificial Intelligence and Robotics. Englewood Cliffs, N.J.: Prentice-Hall, 1985.

Gomaa, Hassan. "The Impact of Rapid Prototyping on Specifying User Requirements." ACM Sigsoft Software Engineering Notes 8, no.2(April 1983):17-28.

Gutis, Philip S. "More Trips Start at a Home Computer." New York Times 23 December 1989, 54.

Hayes-Roth, Barbara, and Frederick Hayes-Roth. "A Cognitive Model of Planning." Cognitive Science 3, no.4(October-December 1979):275-310.

Hekmatpour, Sharam. "Experience with Evolutionary Prototyping in a Large Software Project." ACM Sigsoft Software Engineering Notes 12, no.1(January 1987):38-41.

Henderson, Rebecca M., and Kim B. Clark. "Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms." Administrative Science Quarterly 35, no.1(March 1990):9-30.

Klausner, A., and T. E. Konchan. "Rapid Prototyping and Requirements Specification Using PDS." ACM Sigsoft Software Engineering Notes 7, no.5(December 1982):95-105.

Kogut, Bruce, and Uto Zander. "Knowledge of the Firm and the International Replication of Technology." Wharton School Paper, July 1989.

Larkin, Jill, John McDermott, Dorothea P. Simon, and Herbert A. Simon. "Expert and Novice Performance in Solving Physics Problems." Science 208, no.4450 (20 June 1980):1335-42.

Lave, Jean, Michael Murtaugh, and Olivia de la Rocha. "The Dialectic of Arithmetic in Grocery Shopping." In Everyday Cognition: Its Development in Social Context, edited by Barbara Rogoff and Jean Lave, 67-94. Cambridge, Mass.: Harvard University Press, 1984.

Luchins, A. S. "Mechanization in Problem-Solving: The Effect of Einstellung." Psychological Monographs 54 (1942).

Nelson, Richard R. "Production Sets, Technological Knowledge, and R&D: Fragile and Overworked Constructs for Analysis of Productivity Growth?" American Economic Review: Papers and Proceedings 70, no.2(May 1980):62-67.

Nelson, Richard R. "The Role of Knowledge in R&D Efficiency." Quarterly Journal of Economics 97, no.3(August 1982):453-70

Newell, Allen, and Herbert A. Simon. Human Problem Solving. Englewood Cliffs, N.J.: Prentice-Hall, 1972.

Polanyi, Michael. Personal Knowledge: Towards a Post-Critical Philosophy. Chicago: University of Chicago Press, 1958.

Rogers, Everett M. Diffusion of Innovations. 3rd ed. New York: Free Press, 1983.

Rothwell, R., C. Freeman, et al. "SAPPHO Updated - Project SAPPHO Phase II." Research Policy 3 (1974):278-91.

Scribner, Sylvia. "Thinking in Action: Some Characteristics of Practical Thought." In Practical Intelligence: Nature and Origins of Competence in the Everyday World, edited by Robert J. Sternberg and Richard K. Wagner, 13-30. Cambridge: Cambridge University Press, 1986.

Scribner, Sylvia. "Studying Working Intelligence." In Everyday Cognition: Its Development in Social Context, edited by Barbara Rogoff and Jean Lave, 9-40. Cambridge, Mass.: Harvard University Press, 1984.

Shaw, Brian. "Appropriation and Transfer of Innovation Benefit in the U.K. Medical Equipment Industry." Technovation 4 (1986):45-65.

Shocker, Allan D., and V. Srinivasan. "Multiattribute Approaches for Product Concept Evaluation and Generation: A Critical Review." Journal of Marketing Research 16, no.2 (May 1979):159-80.

Silk, Alvin J., and Glen L. Urban. "Pre-Test-Market Evaluation of New Packaged Goods: A Model and Measurement Methodology." Journal of Marketing Research 15, no.2 (May 1978):171-91.

Simon, Herbert A., et al. "Decision Making and Problem Solving." Interfaces 17, no.5(September-October 1987):11-31.

Teece, David J. "The Multinational Enterprise: Market Failure and Market Power Considerations." Sloan Management Review Spring 1981:3-17.

Tyre, Marcie J. "Managing Innovation in the Manufacturing Environment: Creating Forums for Change on the Factory Floor." Working Paper No. 3005-89-BPS, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Mass., March 1989.

Urban, Glen L., and John R. Hauser. Design and Marketing of New Products. Englewood Cliffs, N.J.: Prentice-Hall, 1980.

von Hippel, Eric. The Sources of Innovation. New York: Oxford University Press, 1988.

von Hippel, Eric. "Successful Industrial Products from Customer Ideas." Journal of Marketing 42, no.1(January 1978):39-49.

von Hippel, Eric. "Task Partitioning: An Innovation Process Variable." Research Policy 19, no.5(October 1990)(forthcoming). Also available in Working Paper No. 2030-88, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Mass., June 1988.

Witte, Eberhard. "Field Research on Complex Decision-Making Processes -- The Phase Theorem." International Studies of Management & Organization 2, no.2 (Summer 1972):156-82.

Zelkowitz, Marvin V. "A Case Study in Rapid Prototyping." Software--Practice and Experience 10, no.12 (December 1980): 1037-42.

Zuboff, Shoshana. In the Age of the Smart Machine: The Future of Work and Power. New York: Basic Books, 1988.