"Sticky Information" and New Marketing Research Methods

Eric von Hippel

WP# 3753-95 January, 1995

ABSTRACT

To enable the design of new products and services that are accurately responsive to user needs, market researchers must bring relevant need information and solution information together with problem-solvers at a common location. Traditionally, market researchers have addressed this problem by designing methods to collect need-related information from users, analyze it and transfer it to manufacturer-based product and service developers. These methods have been successful under many conditions. However, they inevitably encounter difficulties when information related to user needs is "sticky" - very costly or impossible to transfer from users to manufacturers.

In this paper we first show that information related to user needs is in fact often sticky. We then show that, when this is the case, it can be advantageous to shift product or service design activities to the site of sticky information rather than striving to shift the sticky information from users to manufacturer-based developers. Achieving this new partitioning of innovation process tasks will require new marketing research methods that are complementary to existing ones. Extant work towards this end is summarized, and additional possibilities are discussed.

"Sticky Information" and New Market Research Methods

1. Introduction

A central goal of the field of marketing research is to enable the development of products and services that are accurately responsive to user needs. Traditionally, market researchers have addressed this goal by developing methods that collect information related to product and service needs from users and transfer it to manufacturers for analysis and response. Such "manufacturer-based" methods are very effective under many conditions. However, they inevitably encounter difficulties when information related to user needs is "sticky," or costly to transfer from one site to another.

In this paper we explore the relationship between sticky information and types of marketing research methodologies appropriate for new product and service development. We begin by reporting that need and/or solution information needed by product and service designers is often sticky at user and/or manufacturer sites (section 2). Next, we find that patterns in information stickiness can affect the most cost-effective location for product and service design activities. For example, when market researchers and developers require access to very sticky need-related information that is located at user sites, it can be cost-effective to shift product and service design activities to those user sites (section 3). Finally, we explore work being carried forward by a number of researchers that is leading both towards marketing research methods appropriate to user-based design activities, and towards methods appropriate to a pattern in which design activities iterate between user and manufacturer (section 4). These methods will complement existing, manufacturer-based marketing research approaches in that they are appropriate to different points in the range of information stickiness conditions faced by market researchers and developers.

2. "Sticky" Information

The "stickiness" of a given unit of information in any given instance is defined as the incremental expenditure required to transfer that unit of information to a specified site in a form usable by a given information seeker (von Hippel 1994). When this cost is low, information stickiness is low; when it is high, stickiness is high.

The stickiness of a given unit of information is affected by attributes of the information itself, and also by attributes of and choices made by information seekers and information providers. For example, if a particular information seeker is inefficient or less able in acquiring information unit x , (e.g., because of a lack of certain tools or complementary information) or if a particular information provider decides to charge for access to unit x, the stickiness of that unit of information will be higher than it might be under other conditions. The stickiness of a given unit of information is not immutable, and can be affected by investment to that end. For example, firms can invest in encoding some types of skills held by expert employees into the more easily transferrable form of specialized software "expert systems"(Davis 1986) able to partially emulate those skills.

It has not always been clear that information might be significantly sticky. Indeed, the central tendency in economic theorizing has been to view information as costlessly transferable, and much of the research on the special character of markets for information has been based precisely on this characteristic (e.g., Arrow 1962, 614-15). However, more recently it has become clear that the costs of information transfer vary significantly and can be quite substantial (Cohen and Levinthal 1990, Griliches 1957, Mansfield 1968, Nelson 1982 &1990, Pavitt 1987, Rosenberg 1982, Teece 1977). For example, Teece, in a study of 26 international technology transfer projects, found that the costs associated with transferring nonembodied technical knowledge to a new site ranged from 2 percent to 59 percent, and averaged 19 percent of total project costs - a very considerable fraction (Teece 1977, 245, 247).

A full discussion of sticky information has been provided elsewhere (von Hippel 1994) and will not be repeated here. Instead, we will illustrate the phenomenon and its relevance to the concerns of marketing research by reviewing and illustrating two common causes of high information transfer costs - the nature of information encoding and the amount of information that must be transferred during the course of product and service development.

Information Encoding

It is well understood that some information held by users or manufacturers is encoded in explicit terms, while other information is "tacit." Polanyi points out that many human skills and much human expertise are tacit, and illustrates the point by noting that "the aim of a skilful 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, italicized in original). For example, swimmers are probably not aware of the rules they employ to keep afloat, e.g., in exhaling, they do not completely empty their lungs, nor are medical experts generally aware of the rules they follow in order to reach a diagnosis of various symptoms. "Indeed," he says, "even in modern industries the indefinable knowledge is still an essential part of technology." And, Polanyi reasons, "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..." - a relatively costly mode of transfer (ibid., 52,53).

Tacit information is important to market researchers and designers because both the use and the development of products and services generally involves human expertise or skill. For example, manufacturer-based designers who wish to develop a flotation aid for swimmers will be aided by knowing swimmers' tacit rules, so that their design can work in harmony with those rules. Similarly, a user - a swimmer - who wishes to design a flotation aid for swimmers will be aided by knowing some of the manufacturer-based designers' explicit and tacit information regarding practical and producible product design.

The richness of tacit information, and the difficulty of transferring it from person to person or place to place (e.g, from user to manufacturer or vice versa) is nicely illustrated in a study by Barley and Bechky (1994) which documents efforts by some cell culture technicians to transfer their partially tacit skills to others:

> "Talk in both [cell culture] labs routinely centered on the relevance of colors, shapes and patterns, and occasionally on sounds and smells. For instance, the MAb lab's research support specialist constantly referred to the importance of "keeping the cells happy," an idiom she shared with other cell culture specialists. ... To ensure healthy cells, cell culture specialists continually monitored differences in the cells' shape and color as well as changes in the visible properties of the media in which they grew. A sense for the semiotic nature of the work can be gleaned from notes taken as the MAb lab's research support specialist (Sally) trained an inexperienced technician (Mary) to evaluate a [cell] fusion. The discourse turned on calling the technician's attention to the meaning of visible cues:
>
> [Sally and Mary took turns peering through a microscope. Sally looked first. Then, as Mary looked, Sally told her what to notice.] Sally asked Mary to look at the cells. ... Sally called Mary's attention to a well in which the *medium has turned dark yellow*, a sign that the well contained many hybridomas. As Mary examined the well, Sally noted that the *"stretched out"* cells on the bottom of the well were fibroblasts. Mary asked how you could tell if a cell was a hybrid. Sally responded that if it weren't it would eventually die, and that dead cells looked *"dark and grainy."*
>
> * * *
>
> "Although instrument manuals and textbooks discussed many of the signs crucial to the practices of both labs, the staff claimed that only by experience could one become an accomplished interpreter. Experience was deemed critical for two reasons. First, considerable information was

> carried by subtle differences in shading and pattern. Second, like spoken languages, technological codes exhibited "dialects" or local variations. These variations were often tied to peculiarities of specific cell lines, machines, and experiments. Hence, experienced research support specialists and technicians made use of signs that could not be found in textbooks, and that were difficult to define except ostensively. Partially for this reason, practices successful in one lab often failed in another unless technicians from the first trained technicians from the second (for similar observations see Cambrosio and Keating, 1988; Jordan and Lynch, 1992) . (Barley and Bechky 1994 p. 98-9)

This case describes difficulties in the transfer of relatively technical information and skills. However, consider that more common skills such as those associated with making a cake ("fold the egg whites until they appear slightly stiff") or conducting a conversation via telephone ("he hesitated very slightly before answering - I think he may disagree with my suggestion") are also complex and partially held in tacit form by those possessing them. Therefore, we may expect that some key information related to these more common skills will also be difficult to encode and transfer economically to others.

Amount of Information

The cost of transferring information called for by a problem-solver from one location to another can be very high even when the needed information has a low stickiness per "unit" - simply because a great deal of such information may be needed by product and service designers. To understand why this is so, consider that a user's need and use environment and/or the solution skills and solution environment of manufacturer-based designers can contain a myriad of highly specific attributes. When this is the case, it will not be cost-effective to simply transfer all information related to one of these environments to the other. Nor will it be possible to identify and transfer to designers only that subset of information that they will find relevant during their problem-solving work. This is because the relevant subset of information is contingent upon the solution path designers take during their problem-solving work. And, since the problems designers work on are "ill-structured" ones[1], the path and outcome of their problem-solving work cannot be predicted in advance. As illustration of both points, consider a case drawn from von Hippel and Tyre (1994), which involves difficulties associated with transferring need-related information from a user environment to manufacturer-based designers.

Problem of the Yellow Circuit Board

> A process machine development group was asked by a computer manufacturer to design a "component placing machine" to be used in the assembly of printed circuit boards. The machine's function would be to automate the task of placing large integrated circuits at precise locations on those boards. Process machine designers collected information on needs and the use environment from the potential users of the component placing machine, and then began development. As development proceeded, the designers created a machine that was made up of two major components: a "machine vision system" that was used to determine the proper location for each integrated circuit being placed on a circuit board being processed, and

---

[1]Well structured problems are defined as those for which one can precisely specify a process of trial and error that will lead to a desired solution in a practical amount of time (Reitman 1965, Simon 1973, Pople 1982). For example, a traveling salesman problem can be well structured, because one can precisely specify a generator of alternative solutions and a solution testing procedure that are guaranteed to eventually identify the best solution. However, "In general, the problems presented to problem solvers by the world are best regarded as ill structured problems. They become well structured problems only in the process of being prepared for the problem-solvers. It is not exaggerating much to say that there are no well structured problems, only ill structured problems that have been formalized for problem-solvers." (Simon 1973 p. 186).

Ill structured problems may involve an unknown "solution space" (a precisely specifiable domain(s) in which the solution is known to lie). They may also involve unknown or uncertain alternative solution pathways, inexact or unknown connections between means and ends and/or other difficulties. Ill-structured problems are solved by a process of first generating one or more (typically several) alternative solutions. These may or may not be the best possible solutions - one has no way of knowing. These alternatives are then tested against a whole array of requirements and constraints (Marples 1961, Simon 1981 p.149). Test outcomes are used to revise and refine the solutions under development, and - generally - progress is made in this way towards an acceptable result.

> a robot arm and hand that physically picked up the integrated circuits and placed them at those locations.
>
> The input to the machine vision system was a small video camera used to search for particular metalized patterns on the surface of boards being processed. In order for the vision system to function properly, it was necessary that the video camera be able to "see" these metalized patterns clearly against the background color of the board surface itself. The vision system functioned properly in the lab when tested with sample boards from the user plant. However, when it was introduced into the factory, it sometimes failed. Development engineers came to the field to investigate, and found that the failures were occurring when boards that were light yellow in color were being processed.
>
> The fact that boards being processed *were* sometimes light yellow was a surprise to the machine developers. Factory personnel knew that the boards they processed varied in color, but had not volunteered this information because they did not know that the developers would be interested. Early in the machine development process, they had simply provided samples of boards used in the factory to the machine development group. And, as it happened, these samples were green in color. On the basis of the samples, developers had then (implicitly) assumed that all boards processed in the field were green. It had not occurred to them to ask users, "how much variation in board color do you generally experience?" Thus, they had designed the vision system to work successfully with boards that were green.

In retrospect, one can say that the product (a process machine in this case) being developed failed to meet user expectations because an element of information about the use environment had not been transferred from user to manufacturer. Why had it not been? After all, the information on board color variations was known to the users. To understand the difficulties that can attend transferring the "obvious," consider first that the aspect of the use environment at issue in the yellow board case was a very narrow and specific one. That is, the problem with the board was not that it had "physical properties," nor that it had a color. The problem was precisely that the boards were yellow, and a particular shade of yellow at that. Since a circuit board - indeed, most components - have many attributes in addition to color (shape, size, weight, chemical composition, resonant frequency, dielectric constant, flexibility, etc., etc.) it is likely that market researchers seeking to collect all of the information about the user and use environment required to design the product in a way that was precisely responsive to the need would have to analyze a very large (perhaps unfeasibly large) number of potentially problematic items and interactions to achieve this.

Note next that the problem caused by the yellow color of the board was contingent on the particular solution to the component placing problem selected and developed by the engineer, and this was only done during the problem-solving work of engineering design. That is, the color of printed circuit boards in the user factory became relevant only when engineers, during the course of their development of the component placer, decided to use a vision system in the component-placing machine they were designing; the fact that the boards were yellow only became relevant when the engineers chose a video camera and lighting that could not distinguish the metalized patterns on the board against a yellow background. Since engineers often change the alternatives they are developing during the course of their development work (Marples 1961, Allen 1966), it follows that the relevance to designers of any particular item of information bearing on product or service needs - or potential solutions to those needs - can also change frequently during the development process.

In sum, then, we see that some of the need and/or solution information that developers require is likely to be sticky in many product and service development projects.

3. Sticky Information and the Locus of Product and Service Design Activities

Creation of new products or services that are accurately responsive to user need requires that designers have access to both need information and solution information. In the case of product and service design, two information bases located - at least, initially - in physically different places, are typically important for successful problem solving. The first is information on need, located initially with the user. The second is information on solution technologies, located initially at the site of the manufacturer.

Consider three patterns for bringing need information and solution information together with problem-solvers at a common location (Figure 1). Each involves a different subdivision of innovation-related tasks between user and manufacturer, and involves the transfer of different information between these two parties as well.

Manufacturer-based design (Figure 1a) begins with the transfer of need information from users to manufacturers by marketing research and/or sales personnel. Development is then carried out by manufacturer-based designers who draw on the need information that has been transferred to them plus solution information already located at the manufacturer site to create a new product or service that is responsive to user need. Manufacturer-based design is the conventional wisdom as to how product and service development should proceed, and the pattern is doubtless familiar to every reader. (Current marketing research methods are designed to support the manufacturer-based design partitioning of innovation-related tasks.)

(1a) Manufacturer-Based Design

Manufacturer-based design tasks

Have Solution Information
Acquire User Need Information from User
Combine to Design Responsive Product

User-based design task

Have User Need Information

(1b) User-Based Design

Manufacturer-based design task

Have Solution Information

User-based design tasks

Have User Need Information
Acquire Solution Information from Mfr
Combine to Design Responsive Product

(1c) Iterative User and Manufacturer-Based Design

Manufacturer-based activity	User-based activity
	user draws on local need information to specify desired product or service
manufacturer draws on local capability information to develop prototype responsive to specifications	
	user evaluates prototype, drawing on local information regarding application context, and improves/ changes specifications as evidence dictates
manufacturer iterates until user satisfied	
	user iterates until satisfied

Figure 1: Three alternative arrangements of product and service design tasks. Which will be preferred in a given instance depends on the relative stickiness of user-based and manufacturer-based information in that case. (When both need and solution-related information are not sticky, any of the patterns can be used with equal economy from the point of view of information transfer costs.)

User-based design (Figure 1b) involves a user developing a new product or service to satisfy his or her own need. The process begins with the transfer of solution information to users. User-based designers then combine that solution information with their own internal need data, and design responsive new products or services for themselves. User-based designers acquire the solution information they need by drawing on solution "kits" supplied by a manufacturer for that purpose or by assembling the solution information and components and tools they need from a range of manufacturers and other sources. As an illustration of the former approach, consider Object Oriented Programming languages (OOPs). Manufacturers of OOPs basically are in the business of creating and selling "user-friendly" software development tool kits. User-based designers acquire these kits and use them to design software products for themselves that are responsive to their own needs. As an illustration of the latter approach, consider the design activities of production engineers who select and assemble equipment supplied by manufacturers into a customized production system. Similarly, consider the design activities of individual consumers who select and purchase manufacturer-supplied components such as telephone answering machines and note pads and computers and desks and chairs and integrate these into a home office workplace of their own devising.

Finally, in the case of iterative user and manufacturer-based design (Figure 1c), neither sticky need information nor sticky solution information is transferred from its original site. Instead, problem-solving activity is shifted iteratively between user(s) and manufacturer(s) as sticky information located at these two sites is drawn upon by problem solvers. Iteration between the two sites occurs - rather than a one-time access to each site - because problem solving in general (Barron 1988, 43-47) and technical problem solving in particular (Marples 1961, Allen 1966) has trial and error as a prominent feature. If and as each cycle of a trial and error process requires access to sticky information located at more than one site, iterative shiftings of problem solving activity among sticky information sites will occur as problem solving proceeds.

The "rapid prototyping" method used today by some software developers offers a good illustration of iterative user and manufacturer-based design. In that method, manufacturers respond to initially perceived needs for new software by quickly developing and delivering to users an inexpensive, easy to modify, working model that simulates key aspects of the functionality of the planned new software. Users then use the prototype in their own setting on their own data and clarify and modify their perceptions of their needs as a result of this experience. They then relay requests for change or new features to the software developers, who respond by drawing on their own sticky information and tools to make modifications to the prototype. A revised prototype is then sent to the user, and this process of iteration between developer and user is repeated until an acceptable fit between need and solution is found.

We hypothesize that, for any given design project, those with the power to do so will select the pattern of innovation process activities that minimizes their information transfer costs, other things being equal. If we assume that a user and a manufacturer with information relevant to a given project will try to minimize their joint information transfer costs, we reason that manufacturer-based design (Figure 1a) will be preferred for projects where solution information is sticky at the manufacturer's site and need-related information is not sticky; that user-based design (Figure 1b) will be preferred for projects that involve sticky need information and non-sticky solution-related information; and, finally, that iterative user and manufacturer-based design (Figure 1c) will be preferred for design projects involving sticky need information and sticky solution information.

<u>A Model</u>

A simple schematic model can be used to compare the relative efficiencies of user-based and manufacturer-based design with respect to information transfer costs. Assume that the efficiency of user vs manufacturer based design in any particular instance will vary as a function of: (1) The amount (I_u) and the stickiness (S_u) of user-based need information that must be transferred from user to manufacturer if the problem is to be solved via manufacturer-based design; (2) the amount (I_m) and the stickiness (S_m) of manufacturer-based solution information that must be transferred from manufacturer to user if the problem is to be solved via user-based design; (3) the cost of design activities when carried out at the user (D_u) or the manufacturer (D_m) location. The relative efficiency of user-based and manufacturer-based design in a particular instance can then be represented as:

$$\frac{\text{User-based design costs}}{\text{Mfr-based design costs}} = \frac{(I_m S_m + D_u)}{(I_u S_u + D_m)}$$

The relative efficiency of iterative user and manufacturer-based design for a given project can also be assessed using this general model. To do so one would periodically recalculate model values as problem-solving work progressed, in order to determine the points at which shifting between user-based and manufacturer-based design activities would be effective.

Note that the information addressed by our model is not *all* need-related or solution-related information required by project designers, but rather only to that portion of the required information that is *not already in the designer's possession,* and so must be transferred to him or her during the course of the project. To illustrate this distinction, consider a design project to create a cake mix having a new flavor. New cake mixes are often developed using a manufacturer-based design process. Yet, the user need-related information relevant to cake mix development is clearly very rich and complex, and arguably involves a significant amount of sticky information that must somehow be transferred from users to manufacturer-based designers.

For example, to design an acceptable cake mix, designers must understand in detail what a cake is and what it should look like; understand the role it plays in meals and social occasions; understand how it is eaten and so forth. They must also understand what users expect a "cake mix" to be, understand the nature of the baking skills users possess, the nature of the kitchen equipment commonly available to users, etc.. Despite this need for complex and sticky need-related information, manufacturer-based designers may have no difficulty creating a successful cake mix. Both designers and users share a rich cultural context which includes cakes and cake mixes, and so the designers already <u>know</u> the great bulk of the need-related information they will require. Under these conditions, therefore, the information *which is not already in the designer's possession* and so must be transferred from user to designer may be relatively limited and non-sticky - consisting, for example, only of user perceptions and preferences regarding the proposed new cake mix flavor.

The variables in our model, information amount (I) and information stickiness (S), will be usually difficult to specify and measure under real-world conditions for a number of reasons. For example, consider that designers may change a problem as they work on it - and as a result change the nature and amount of information they may require to solve it. Further, recall that information stickiness is determined by the attributes of information providers and users as well as on attributes of the information itself. Nonetheless, the model can serve as a useful conceptual tool to help us think about and perhaps even "guesstimate" the relative cost-effectiveness of user-based vs manufacturer-based design in cases or circumstances of interest. Consider the following two cases as illustration.

Suppose, first, that the product design task in question is the creation of a new generation of computer memory chips (DRAMs). We may speculate that the *novel* user need information required by the DRAM designers in that case will be relatively small in amount and relatively non-sticky - essentially consisting, let us say, of "make it function like the last generation - but faster and cheaper, please." In contrast, it is reasonable that the solution information required to create a new generation of DRAM chips will be very rich and complex and, since DRAMs are at the frontier of the chipmaking art, will be held partially in tacit form in the minds of talented design and process engineers. Thus, we reason that $I_m S_m > I_u S_u$ in this instance. Next, with respect to design costs, note that the design of a new generation of DRAMs requires experiments carried out on complex and costly laboratory equipment. DRAM manufacturers already have much of the needed equipment and related expertise in use and available for a DRAM design project: DRAM users do not. On this basis we may speculate that $D_m < D_u$ in the case of DRAM design. Placing these guesstimated values in our model leads us to conclude that DRAM design will be more economically carried out via manufacturer-based rather than via user-based design.

As a second example, suppose that the design task in question is the creation of a curved-wire probe with a loop at the end to be used by surgeons as a tool to aid in the removal of plaque from the walls of certain arteries during heart operations. Here we reason that the relevant solution information (consisting, let us say, of how one bends and forms the grade of wire used in such surgical tools) is relatively small in amount and non-sticky. In contrast, it is likely that the user-based need information is both voluminous and largely tacit, consisting of the complex interactions between the surgeons' skill and the characteristics of patients' bodies and arteries and of plaque. Thus, we conclude that $I_mS_m < I_uS_u$ in this case. Further, testing and adjusting proposed designs for the probe involves problem-solving under real or very realistic operating theatre conditions. Such conditions are routinely available to practicing heart surgeon users, but not to instrument manufacturers. Therefore it is likely that ($D_m > D_u$). Placing these values in our model leads us to conclude that this surgical instrument design task will be most economically carried out via user-based design rather than via manufacturer-based design.[2]

---

[2] Note that our schematic model addresses the relative efficiency of user-based vs manufacturer-based design from the point of view of an individual user or manufacturer only. If we wish to consider the matter from the point of view of an entire market, we must add information on matters such as the relative number of users and manufacturers participating in that market, the power and incentives they may have with respect to forcing other parties to bear any information transfer costs, the relative efficiency with which user-based versus manufacturer-based designs diffuse once they are developed and so forth.

For example, if both user-developed designs and manufacturer-developed designs diffuse equally well or poorly, then the equation given earlier holds independent of market structure considerations. If diffusion efficiency differs, however, the relative number of users and manufacturers interested in a given design does affect market-level efficiencies for user-based vs manufacturer-based design. Suppose, for example, that several or many users (n) have an identical need for a given new product or service. Suppose further that users who design a solution would absolutely refuse to share their design information, while manufacturers who undertake manufacturer-based design would share that information freely by selling the product or service embodying it. In that case, the equation given earlier would change to:

$$\frac{\text{User-based design costs}}{\text{Mfr-based design costs}} = \frac{(I_mS_m + D_u)}{(I_uS_u + D_m)/n}$$

on a per user basis. In other words, under such conditions the relative efficiency of manufacturer-based design at the level of a market would increase as the number of potential users of a given product or service design increases In fact, however, there may well be no major difference in the rate of diffusion of design information developed by users or by manufacturers. It has been found that detailed information on user-developed processes and process machinery diffuses to rival users in a matter of months (Mansfield 1985). Further, "lead user" market research studies have shown users generally willing to share their design information with inquiring manufacturers (Urban and von Hippel 1988, Herstatt and von Hippel 1990). Once user design information has been transferred to even a single manufacturer, further diffusion can occur via the same route as that taken by manufacturer-developed design information - by the sale of products or services that embody it.

Some Evidence

Our model is based on the assumption that problem-solving activities will be located so as to minimize joint user and manufacturer information transfer costs, other things being equal. This assumption certainly requires test and refinement, but we can note that extant research findings are compatible with it. Mowery and Rosenberg (1989, chap.4) examined the content of all projects carried out by three major independent R&D contracting firms (the Mellon Institute, the Battelle Memorial Institute and Arthur D. Little, Inc.) from 1900 and 1940. They found that the bulk of the projects carried out by the independent R&D contractors were of a nature that required a relatively small amount of firm-specific knowledge, and argued that the projects requiring large amounts of such knowledge had been carried out in client firms' internal labs due to the difficulties of transferring it to another site.

In a similar vein, Riggs and von Hippel (1994), in a study of 64 major improvements affecting two types of scientific instruments, found that 82% of the innovations that allowed users to do qualitatively new types of things were developed by users. In sharp contrast, 87% of innovations that increased instrument convenience and reliability were developed by manufacturers, as were 52% of the innovations that improved the instruments along known dimensions of merit such as sensitivity and resolution. Information stickiness was not measured in this study, but it seems likely that the need-related information required by problem-solvers was stickiest for innovations in the first category. If so, the preponderance of innovations by users in this category - that is, innovations developed by problem-solving activities carried out at user sites - is what we would expect given efforts to minimize information transfer costs.

Type of improvement provided by innovation	Innovation developed by: %User	%Mfr	(n)
(1) New functional capability	82%	18%	17
(2) Improvement to convenience or reliability	13%	87%	24
(3) Improvement to sensitivity, resolution or accuracy	48%	52%	23
		Total	64

Table 1: Source of innovations by nature of improvement effected (Riggs and von Hippel (1994 table 3)

There is also some empirical evidence bearing on our hypothesis that iterative shiftings of problem-solving activity between user and manufacturer sites may occur when both sites hold sticky information. Tyre and von Hippel (1993, Figure 2) explored the innovation-related problem solving involved in identifying and diagnosing 27 field failures in novel process equipment used to automatically assemble complex circuit boards. They observed repeated shifts in the locus of technical problem-solving activity occurring during this work, as problem-solvers iteratively sought access to sticky information located at the machine development laboratory or the factories where the novel machine was being used. The number of shifts found ranged from 0 to 7, and averaged about 2.3 times per problem identified and diagnosed (Figure 2).

FIGURE 2

NUMBER OF SHIFTS BETWEEN PLANT AND LAB DURING PROBLEM SOLVING

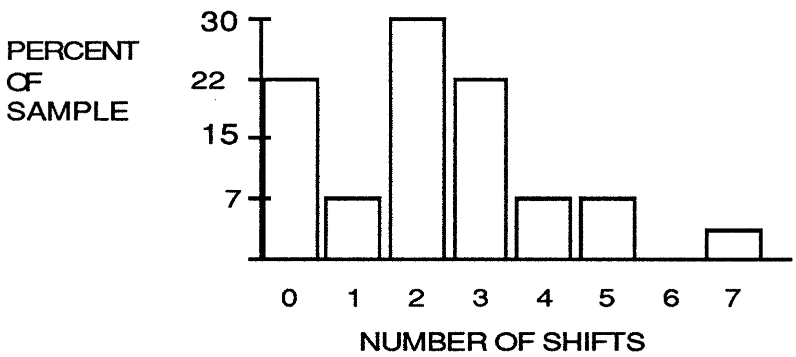

Similarly, Kristensen (1992) found iterative shifts in the location of problem-solving activities in many product development projects carried out by 103 Danish food manufacturers - shifts driven, it appears to us, by the need to gain access to sticky information located at different sites. Thus, he describes the process used by a Danish bakery firm developing a new frozen unbaked cake for a British retail food chain. In this case the producer responded to an initial request from the British customer by developing several prototypes of the proposed cake and sending them to the customer to be baked, tasted, and smelled, and evaluated on the basis of "local tastes and the situation they were meant for - a type of social gathering not practiced in Denmark." Comments on the baked cakes were then sent back to the producer, who adjusted the recipes accordingly, "using his familarity with baking and with local raw materials." In total, "five successive revised generations were sent during the course of three months before the Danish producer and the UK retail chain's test kitchen reached the generation of satisfactory variations." (Kristensen 1992, 204-5, 210).[3]

[3] Note that iterative shiftings of design-related problem solving between user and manufacturer will be less costly than the transfer of sticky information to a single problem-solving locus given a key condition: The intermediate outputs of problem solving conducted at each locus that are transferred between sites must be less sticky than the information operated upon to produce

4. Sticky Information and New Marketing Research Methods

Present-day marketing research methods, ranging from multiattribute analysis (Lancaster 1971, Silk and Urban 1978) to QFD (Hauser and Clausing 1988) to concept engineering (Burchill 1992, Shiba et. al. 1993) are designed to collect need-related information from users and transfer it to manufacturer-based personnel for analysis and responsive action. Such methods seek to implement the "manufacturer-based design" partitioning of product and service design tasks. As we have seen, manufacturer-based design can be cost-effective when the user need information required by designers is relatively unsticky, but must encounter difficulties when information related to user needs is sticky. Clearly, new market research methods that are effective under the latter conditions would be useful. In this section we review method developments on the part of a several researchers that seem to us to be addressed to this goal, and suggest further opportunities as well.

Marketing Research Methods and Iterative User and Manufacturer-Based Design

In some fields, marketing research methods have been developed that specifically incorporate user-manufacturer iteration. We think further effort in this direction would be useful. We also propose that it will be useful to develop methods for converting design problems that draw on sticky information located at two sites into subproblems that each draw on only one site of sticky

---

the outputs. Intuitively it seems reasonable that this will often be the case: Such an intermediate output may be in the form of nonembodied information transferable at low cost, or it may be in the form of a prototype that can be economically transferred. For example, an artist may not be able to transfer all information involved in the creative process that brings him or her to specify to a supplier, "I need a green paint of precisely X hue and luminance." However, that (nonembodied) need specification is very simple and precise, and it can be transferred at very low cost. Similarly, the responding paint manufacturer may be able to create and transfer the requested shade of green to the artist (embodied in a prototype or final product), but not be able to transfer the complex knowledge drawn on by that firm's chemists to achieve the feat.

information. We will argue that such conversions can be very useful to manufacturers and users of many products and services. They are currently being done on an ad hoc basis in some fields but, to our knowledge, marketing research methods to clarify and routinize the procedure have not yet been developed.

With respect to the first proposal, recall that iterative relocation of design activities between user and manufacturer is indicated when both need information and solution information required by product and service designers is sticky. Under these conditions traditional methods will not do, because one cannot expect to fully and accurately transfer need-related information to manufacturer-based designers. As illustration, consider the traditional, specification-driven ("waterfall") method of software development. In that traditional method, systems analysts based at the software developer meet with users at the start of a project to determine user needs and agree on a written product requirements specification. Developers then work isolated from further user contact until the completed product is delivered months or even years later. All too often, software developed according to this method has been "late, over budget and not what the customer wanted" (Zelkowitz 1980, 1037).

Introduction of meaningful iteration into such methods involves breaking manufacturer design activities into a series of steps that each produce an output that can serve as the basis for significant user-based problem-solving. The "rapid prototyping," method of software development mentioned earlier illustrates how this can be done. Rapid prototyping begins with manufacturers responding to initial user need inputs by quickly developing and delivering to users (usually within weeks) an inexpensive, easy to modify, working model that simulates important components of the functionality of the proposed new software. Users then learn by using the prototype in their own setting on their own data and clarify their needs, in part by drawing on their tacit knowledge and experience

(Gronbaek 1989, 114-16). Users then relay requests for change or new features to the software developers, who respond by drawing on their own sticky information and tools to make modifications to the prototype. Some of these modifications are minor, such as altering report formats, and some are major, such as implementing a new feature or modifying the basic structure of the prototype (Feld 1990, 14). A revised prototype is then sent to the user, and this process of iteration between developer and user is repeated until an acceptable fit between need and solution is found.

A number of individual case studies and experiments have shown that rapid prototyping methods "better satisfy true user requirements and produce information and functionality that is more complete, more accurate, and more meaningful [than do traditional, non-iterative methods]" (Connell and Shafer 1989, 15; Boehm, Gray, and Seewaldt 1984; Gomaa 1983) - a result that we would expect when both need and solution information are sticky.

Rapid prototyping methods appear to be appropriate for iterative user and manufacturer-based problem-solving in many fields. When manufacturers create partial or full product and service prototypes, they embody their sticky solution information in a form that can be shifted to user sites. Users at those sites can then apply sticky need-related information to test and improve those prototypes during the development process. Relatively recent technological advances are making rapid prototyping methods steadily more cost-effective by reducing the costs of creating appropriate prototypes. For example, software prototyping costs are being reduced by object oriented programming; mechanical prototyping costs are being reduced by computer aided design and manufacture (CAD-CAM); electronic prototyping costs are being reduced by the introduction of integrated circuits (EPLDs) that users can tailor for their custom application using only a personal computer.

We now turn to our second proposal for research and method development related to iterative user and manufacturer-based design. The "iteration solution" to the problem of needing to draw repeatedly on two or more sites of sticky data during problem-solving can represent the most economical partitioning of a given product or service development project - but it may still be undesirably costly. The reason is that there are costs and time lags inevitably associated with starting up and shutting down problem-solving activities - and patterns such as iterative user and manufacturer-based design incur these costs several times in a single project. Accordingly, if a product or service design problem as originally given requires access to two or more sites of sticky information during the course of problem-solving, it may be desirable to change the original problem by converting it into subproblems that each require access to only one such site. This can be done by "reframing" design problems and/or changing the stickiness of some key information.

As a schematic illustration of reframing, suppose that a pipe manufacturer is given the job of designing a special pipe that must cross a busy construction project, and suppose that this pipe must be manufactured with many precisely-located turns and bends to avoid interfering with sites of present and potential construction activity. Clearly this problem, framed in this way, is appropriate for iterative user and manufacturer-based design. Information about the pathway the pipe must follow is sticky - complex and unpredictably changing as construction proceeds - so transfer of need information to the manufacturer would be quite costly. Nor can the manufacturer easily transfer solution information to the user - the process of designing curves into the pipe is quite sticky, requiring an understanding of shapes that can be cast successfully in a metal foundry. The result will clearly be a lot of iteration between user and manufacturer before the design is gotten right.

Next reframe the original problem into two subproblems: (1) how to design a flexible pipe that can be bent by users at field sites without compromising its ability to function and, (2) design of a pathway across the construction site for the pipe that will not interfere with construction activities. The pipe manufacturer can now address problem (1), design of a flexible pipe, by drawing only on sticky information located at the manufacturer site, while the user can address problem (2), locating the pipe, by drawing only on information located at the user site. Thus, this reframing of the problem eliminates the need to iterate between user and manufacturer during problem-solving in order to gain access to sticky information.

As a real world case of reframing in conjuction with unsticking some key information, consider the problem-solving work involved in designing an integrated circuit for a custom application. In this design problem, two sticky data bases are central to the problem-solving work: (1) information at the circuit user locus involving a rich and complex understanding of both the overall application in which the custom integrated circuit will play a role and the specific function required of that circuit; (2) information at the circuit manufacturer locus involving a rich and complex understanding of the constraints and possibilities of the silicon fabrication process that the manufacturer uses to produce integrated circuits.

Traditionally, custom integrated circuits were developed via an iterative user and manufacturer-based design process involving a circuit user possessing sticky need information and an integrated circuit manufacturer possessing sticky information about designing and producing custom integrated circuits. That process began with a user specifying the functions that the custom chip was to perform to a circuit design specialist employed by the integrated circuit manufacturer. The chip would then be designed at the manufacturer locus, and an (expensive) prototype would be produced and sent to the user. Testing by the user would typically reveal faults in the chip and/or the initial specification, responsive changes would be made, a new prototype built, and so forth.

More recently, the Application Specific Integrated Circuit (ASIC) method of making custom integrated circuits has come into wide practice. In the ASIC method, the overall problem of designing custom circuits has been reframed into two new subproblems which each draw on only one locus of sticky information, thereby eliminating the need to iterate between two such sites in the design process. The manufacturer of ASICs draws on its own sticky information to develop and improve the fabrication processes in its manufacturing plant, a "silicon foundry." The manufacturer also draws on its own sticky information to design "standard" silicon wafers that contain an array of unconnected circuit elements such as logic gates. These standard circuit elements arrays are designed by the manufacturer to be interconnectable into working integrated circuits by the later addition of custom interconnection layers designed in accordance with the needs of specific users.

To facilitate this user task, the manufacturer has invested in unsticking some key information related to the capabilities of its production process, and has encoded it in the form of a user-friendly Computer-Aided Design (CAD) software package. With the aid of this CAD tool, users can design a custom interconnection layer design to meet their specific application needs and yet stay within the production capabilities of the manufacturer's silicon foundry. The software also allows the user to simulate the function of the custom circuit under design, and to conduct trial-and-error experiments. Taken together, these capabilities allow the user to both design a circuit, and to refine need specifications and the desired circuit function through an iterative process that draws only on sticky information located at the user site.

Instances of problem reframing to escape iterative user and manufacturer-based design can be seen in a number of fields. The development of "desk-top publishing," to replace iterative problem-solving between a graphic designer and an author, is one such example. More generally, there is a trend in software (Feld 1990) and other fields towards "empowering users" by reframing products and service design problems in such a way as to create the possibility for user-based design activities that can be conducted independent of the manufacturer. (For example, software manufacturers might create a line of user-friendly programming "tool boxes." Users would then draw on their own sticky information to create software precisely adapted to their needs.) Such reframings offer a way for manufacturers to seek economies by producing standard products, while at the same time enabling users to carry out the problem solving needed to adapt these to specific local needs and conditions.

Currently, the problem-solving work involved in reframing design problems to deal with sticky information transfer issues is not supported by marketing research methods. We propose that it will be useful to develop such methods, but do not ourselves yet understand the shape they should take. However, we can note that those who wish to develop methods of this kind can draw upon literature that has explored the subdivision of tasks for problem-solving purposes (e.g, Marples 1961, Alexander 1964, Simon 1973, von Hippel 1990), and upon expertise in fields such as software where reframings of the desired kind have successfully achieved.

Marketing Research Methods and User-Based Design

In user-based design, as is shown in figure 1, information regarding needs for new products and services is both generated within and used within user firms. This means that manufacturers that produce products and services created via user-based design can do so without having to understand the user needs they satisfy. (For example, in the ASIC process of integrated circuit design and manufacture that we reviewed above, circuit manufacturers were able to produce circuits that were accurately responsive to user needs without themselves having to understand those needs.) If the traditional role of manufacturer-based marketing research is to understand user needs, what tasks can or should it adopt in the case of user-based design? Some researchers have shown that it will be useful to identify user-developed products or services that may be profitably produced by manufacturers. In addition, we propose that it will be useful to develop methods for sensing when the boundaries appropriate to user-based design are crossed.

With respect to first point, note that users design and fabricate products and services to serve their own needs, and do not necessarily know or care whether others would have an interest in the same design. However, there is evidence that sometimes others will find a user-designed product or service valuable, and that in such cases a user-developed product or service may represent a commercially attractive opportunity for a manufacturer. (The potential commercial value of user-developed innovations has been documented in a number of fields by tracing samples of commercially significant innovations to their origins. Many have been found to have been initially designed, prototyped and applied by users [e.g., von Hippel 1976 & 1988, Shaw 1985]).

Novel marketing research methods will be useful to identify the subset of user designs and innovations that may be commercially attractive from a manufacturer's point of view. One extant method that can serve this purpose seeks to identify "lead users" who are both in advance with respect to a general marketplace need, and who have a high incentive to innovate with respect to that need. The method has been executed successfully in a number of applications (Urban and von Hippel 1988, Herstatt and von Hippel 1990). A variant of this method directly searches for published reports of innovations by users and then screens these for commercial attractiveness from a manufacturer's point of view (Bailetti and Guild 1991). (Once a promising user innovation has been identified, it is usually not difficult to transfer the information required to reproduce it from user to manufacturer - even though the need-related information relevant to that innovation might have been too sticky to transfer. This is because information related to user designs is typically encoded in physical prototypes in the case of products, and in explicitly specified routines and behaviors in the case of services.)

With respect to the second point, note that user-based designers incorporate manufacturer-produced products and services as components in their designs and/or use them as tools to execute those designs. Thus, user-based designers of custom integrated circuits require manufacturer-supplied CAD tools and wafers to design and fabricate those custom circuits. Similarly, users who design and build their own custom furniture require manufacturer-produced components (such as wood, nails and glue) and tools (pencils, saws, hammers, and glue guns) to execute their designs.

Manufacturer-supplied components and tools have limits with respect to the range of conditions they can address. Thus, an ASIC manufacturer will inform the user-designer that the space on a single custom chip is limited, and that he or she will only be able to fit a circuit of up to x "size" (e.g., up to x thousands of logic gates) on a single chip. Similarly, the manufacturer of nails specifies that a single nail of x size can only support y weight and that it will or will not rust if placed in contact with water. When the user wishes to design within the constraints specified, the manufacturer need not modify the standard tool or component it supplies for that user, and the user can proceed with user-based problem-solving activities independent of the manufacturer. However, when the user wishes to design outside of the specified constraints, the problem he or she is addressing is thereby converted from one appropriate to user-based design to one requiring iterative user and manufacturer-based design or to one including a subproblem addressable via manufacturer-based design.

It would be useful, we think, to develop marketing research methods to sense when users are approaching the limits of a given manufacturer-supplied product, service or tool, and initiate appropriate manufacturer-based responses. Sometimes user-based designers themselves will take the initiative to begin an iterative user and manufacturer design process or a related manufacturer-based design process by contacting manufacturers when they reach such a limit. (For example, a user designer who has a strong need to "put just a few more transistors" on a single chip than silicon foundry rules allow may well call up foundry engineers to "see if we can work something out.") When this is the case market researchers might find it useful to screen such requests looking for common needs and attractive opportunities. In other cases, however, user designers affected by a limitation on a particular manufacturer-supplied component may simply substitute a different component or design approach instead of informing the manufacturer of the problem. Developing marketing research methods capable of initiating appropriate manufacturer responses under these conditions will be more difficult.

Marketing Research Methods and Manufacturer-Based Design

Marketing research methods that implement manufacturer-based design have been developed to a high level of sophistication. The ability of these methods to (1) collect and (2) analyze sticky information can be improved, and some are working to this end.

Today's market research methods typically collect need-related information via individual or group interviews or questionnaires. Both methods clearly have only a limited ability to transfer information that is sticky because it is poorly-encoded (e.g., tacit information) or involves a great deal of detail - and yet we have seen that manufacturer-based designers will often require the transfer of information having these characteristics. A number of investigators are working to improve the effectiveness of procedures applicable to collecting sticky information from users. Thus, Zaltman (1993) has improved interviewers ability to collect non-verbal information during interviews by asking interviewees to bring along and comment upon relevant visual materials. Others are systematizing methods for collecting information during actual visits to customer sites (Holtzblatt and Jones 1990, Shiba et al 1993). Holtzblatt and Jones (1990 p. iii) explain the information transfer benefits of day-long field visits as follows:

> "The contextual inquiry approach is based on field research techniques, and focuses on interviewing users in their own context as they do actual work.... If we just ask customers what they need, they are unable to tell us. Customers are experts in their work, but they usually cannot articulate the key elements of their work. Similarly, if we only observe customers' actions, we might misinterpret the meaning of their actions. ... Whenever we design, we have assumptions about the nature of the customers' work and how technology solves their problems. These assumptions can be blind spots that keep us from seeing information that challenges our assumptions. Contextual inquiry provides a way to align our understanding with customers' understanding. We expand our entering understanding by probing things we do not understand, behavior that surprises us, and problems behind solutions that customers offer. We share our interpretations with customers to create a shared understanding .

After need-related information is collected, it is analyzed by manufacturer-based market researchers. Current quantitative analytical methods tend to strip much of the richness from need-related data that is collected from users, and some are working on methodological improvements that can reduce such losses. We can illustrate the problem by reference to "multiattribute" analytical methods (Lancaster 1971, Silk and Urban 1978, Shocker and Srinivasan 1979, Urban and Hauser 1994) that are frequently used in quantitative marketing research today.

Multiattribute analysis begins when the records of user interviews, site visits, etc. are examined by a market researcher. The researcher's goal is to encode this qualitative information in the form of about 20 scalable, independent component attributes which can be analyzed quantitatively. (The method used by the analyst to identify or create the set of component attributes is typically some formal or informal type of content analysis. The number of attributes identified is limited to about 20 to make succeeding analytical work more tractable.) Once a set of attributes has been specified, the stage for quantitative analysis has been set. A consumer's perception of any particular product in the category can then be expressed quantitatively terms of the amounts of each attribute the consumer perceives it to contain, and the difference between any two products in the category can be expressed as the differences in their attribute profiles. Potential wants and demands for a product or service containing any mix of component attributes can also be determined by including consumer data on the importance and desirability of each of the component product attributes in the analysis.

It will be clear to the reader that quantitative analyses of this type can produce very useful findings. It will also be clear, however, that the analyses are based on only a small portion of the information that collected or collectable from users. Current efforts to improve the richness of information actually transferred to manufacturer-based designers of new products and services tend to involve supplementing information generated by quantitative analyses with qualitative information. Procedures such as QFD (Hauser and Clausing 1988) and Customer Requirements Analyses (Shiba et al 1993) are then used to help designers and marketers to integrate and manage both types of information during the problem-solving work of product and service design.

5. Conclusion

In this paper we have explored opportunities for the development and improvement of marketing research methods that spring from an understanding of the effects of sticky information on patterns and costs of information transfer.

To this point, the great preponderance of method development has been focused on bringing what we have termed "manufacturer-based design" methods to a high level of efficiency. However, analyses of the effects of information stickiness show that it will be useful to enhance marketing research methods related to user-based and iterative user-manufacturer based patterns of design activities as well.

References

Alexander, Christopher (1964), Notes on the Synthesis of Form, Cambridge, MA: Harvard University Press.

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

Arrow, Kenneth J. 1962. "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, ed. Richard R. Nelson, 609-25. Princeton, N.J.: Princeton University Press.

Bailetti, Antonio J. and Paul D. Guild 1991. " A Method for Projects Seeking to Merge Technical Advancements with Potential Markets" R & D Management vol 21 no.4 pp. 291-300.

Barley, Stephen R. and Beth A. Bechky 1994. "In the Backrooms of Science: The Work of Technicians in Science Labs." Work and Occupations, Vol 21 No. 1, Febrauary 1994 85-126.

Barron, Jonathan. 1988. Thinking and Deciding. New York: Cambridge University Press.

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

Burchill, Gary 1992. "Concept Engineering." The Center for Quality Managment, Cambridge Mass.

Cohen, Wesley M., and Daniel A. Levinthal. 1990. "Absorptive Capacity: A New Perspective on Learning and Innovation." Administrative Science Quarterly 35, no.1 (March): 128-52.

Connell, John L., and Linda Brice Shafer. 1989. Structured Rapid Prototyping: An Evolutionary Approach to Software Development. Englewood Cliffs, N.J.: Prentice-Hall.

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

Feld, Bradley A. 1990. "The Changing Role of the User in the Development of Application Software." Working Paper No. BPS 3152-90, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Mass., August 1990.

Griliches, Zvi. 1957. "Hybrid Corn: An Exploration in the Economics of Technical Change." Econometrica 25, no.4 (October):501-22.

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

Gronbaek, Kaj. 1989. "Rapid Prototyping with Fourth Generation systems -- An Empirical Study." Office: Technology and People 5, no.2 (September):105-25.

Hauser, John R., and Don P. Clausing. 1988. "The House of Quality." Harvard Business Review 66, no.3 (May-June): 63-73.

Herstatt, Cornelius, and Eric von Hippel. 1992. "From Experience: Developing New Product Concepts Via the Lead User Method: A Case Study in a "Low Tech" Field", Journal of Product Innovation Management 9: 213-221.

Holtzblatt, Karen and Sandra Jones (1990). "Contextual Inquiry: Principles and Practice." Digital Equipment Corporation Working Paper DEC-TR 729 (October). Digital Equipment Corporation, Maynard, MA.

Katz, Ralph, and Thomas J. Allen. 1982. "Investigating the Not Invented Here (NIH) Syndrome: A Look at the Performance, Tenure, and Communication Patterns of 50 R&D Project Groups." R&D Management 12, no.1 (January):7-19.

Katz, Ralph, and Thomas J. Allen. 1988. "Organizational Issues in the Introduction of New Technologies." In Managing Professionals in Innovative Organizations, ed. Ralph Katz, 442-56. Cambridge, Mass.: Ballinger.

Katz, Ralph, and Michael L. Tushman. 1980. "External Communication and Project Performance: An Investigation into the Role of Gatekeepers." Management Science 26, no 11 (November): 1071-85.

Kristensen, Preben Sander. 1992. "Flying Prototypes: Production Departments' Direct Interaction with External Customers." International Journal of Operations & Production Management 12, no.7,8:195-211.

Lancaster, Kelvin. Consumer Demand.1971. New York: Columbia University Press.

Mansfield, Edwin. 1968. Industrial Research and Technological Innovation: An Econometric Analysis. New York: W.W. Norton.

Mansfield, Edwin. 1985. "How Rapidly Does New Industrial Technology Leak Out?" Journal of Industrial Economics 34, no.2 (December): 217-23.

Marples, David L. 1961. "The Decisions of Engineering Design." IRE Transactions on Engineering Management, June:55-71.

Mowery, David C., and Nathan Rosenberg. 1989. Technology and the Pursuit of Economic Growth. New York: Cambridge University Press.

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

Nelson, Richard R. 1990. "What is Public and What is Private About Technology?" Consortium on Competitiveness and Cooperation Working Paper No. 90-9. Berkeley, Calif.: Center for Research in Management, University of California at Berkeley, April 1990.

Pavitt, Keith. 1987. "The Objectives of Technology Policy." Science and Public Policy 14, no.4 (August): 182-88.

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

Pople, Harry E. Jr. (1982) "Heurisitic Methods for Imposing Structure on Ill-Structured Problems: The Structuring of Medical Diagnostics," Chapter 5 in Peter Szolovits, ed: Artificial Intelligence in Medicine Westview Press, Boulder, Colorado

Reitman, W. R. (1965) Cognition and Thought Wiley, New York

Riggs, William and Eric von Hippel. 1994. "The Impact of Scientific and Commercial Values on the Sources of Scientific Instrument Innovation," Research Policy 23 (July): 459-469.

Rosenberg, Nathan. 1982. Inside the Black Box: Technology and Economics. New York: Cambridge University Press.

Shaw, Brian. 1985. "The Role of the Interaction between the User and the Manufacturer in Medical Equipment Innovation." R&D Management 15, no.4 (October):283-92.

Shiba, Shoji, Alan Graham and David Walden (1993) A New American TQM: Four Practical Revolutions in Management Productivity Press, The Center for Quality Managment, Cambridge Mass.

Shocker, Allan D. and V. Srinivasan (1979). "Multiattribute Approaches for Product Concept Evaluation and Generation: A Critical Review," Journal of Marketing Research 16 (May), 159-80.

Silk, Alvin J. and Glen L. Urban (1978). "Pre-Test-Market Evaluation of New Packaged Goods: A Model and Measurement Methodology," Journal of Marketing Research 15 (May), 189.

Simon, H. A. (1973) "The Structure of Ill Structured Problems," Artificial Intelligence 4, 181-201

Simon, Herbert A.(1981) The Sciences of the Artificial, Second Edition. Cambridge: MIT Press.

Taylor, Mark Peter. 1991. Innovations from User-Initiated Need Messages Unpublished S.M. thesis, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Mass. (June).

Teece, David J. 1977. "Technology Transfer by Multinational Firms: The Resource Cost of Transferring Technological Know-How." Economic Journal 87, no.346 (June): 242-61.

Tyre, Marcie J., and Eric von Hippel. 1993. "Locating Adaptive Learning: The Situated Nature of Adaptive Learning in Organizations." Working Paper No. BPS 3568-93, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Mass., May 1993.

Urban, Glen L., and John R. Hauser. Design and Marketing of New Products. Second Edition, Englewood Cliffs, N.J.: Prentice-Hall, 1992.

Urban, Glen L., and Eric von Hippel. "Lead User Analyses for the Development of New Industrial Products." Management Science 34, no. 5 (May 1988):569-82.

von Hippel, Eric. 1976. "The Dominant Role of Users in the Scientific Instrument Innovation Process," Research Policy 5, no. 3 (July):212-39.

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

von Hippel, Eric. 1990. "Task Partitioning: An Innovation Process Variable." Research Policy 19, no.5 (October): 407-18.

von Hippel, Eric. 1994. "Sticky Information" and the Locus of Problem Solving: Implications for Innovation" Management Science 40, no.4 (April): 429-439

von Hippel, Eric and Marcie Tyre. 1993. "How "Learning by Doing" is Done: Problem Identification in Novel Process Equipment," Working Paper # BPS 3521-93, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Mass., January, 1993 (Research Policy, forthcoming 1994)

Zaltman, Gerald and Robin A. Higie. 1993. "Seeing the Voice of The Customer: The Zaltman Metaphor Elicitation Technique." Working paper #93-114, Marketing Science Institute, Cambridge, MA.

Zelkowitz, Marvin V. 1980. "A Case Study in Rapid Prototyping." Software -- Practice and Experience 10, no.2 (December):1037-42.