


Democratizing Innovation

ERIC VON HIPPEL



© 2005 Eric von Hippel

Exclusive rights to publish and sell this book in print form in English are licensed to The MIT Press. All other rights are reserved by the author. An electronic version of this book is available under a Creative Commons license.

MIT Press books may be purchased at special quantity discounts for business or sales promotional use. For information, please email special_sales@mitpress.mit.edu or write to Special Sales Department, The MIT Press, 5 Cambridge Center, Cambridge, MA 02142.

Set in Stone sans and Stone serif by The MIT Press. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Hippel, Eric von.

Democratizing innovation / Eric von Hippel.

p. cm.

Includes bibliographical references and index.

ISBN 0-262-00274-4

1. Technological innovations—Economic aspects. 2. Diffusion of innovations.

3. Democracy. I. Title.

HC79.T4H558 2005

338'.064—dc22

2004061060

10 9 8 7 6 5 4 3 2 1

Why does a user wanting a custom product sometimes innovate for itself rather than buying from a manufacturer of custom products? There is, after all, a choice—at least it would seem so. However, if a user with the resources and willingness to pay does decide to buy, it may be surprised to discover that it is not so easy to find a manufacturer willing to make exactly what an individual user wants. Of course, we all know that mass manufacturers with businesses built around providing standard products in large numbers will be reluctant to accommodate special requests. Consumers know this too, and few will be so foolish as to contact a major soup producer like Campbell's with a request for a special, "just-right" can of soup. But what about manufacturers that specialize in custom products? Isn't it their business to respond to special requests? To understand which way the innovate-or-buy choice will go, one must consider both transaction costs and information asymmetries specific to users and manufacturers. I will talk mainly about transaction costs in this chapter and mainly about information asymmetries in chapter 5.

I begin this chapter by discussing four specific and significant transaction costs that affect users' innovate-or-buy decisions. Next I review a case study that illustrates these. Then, I use a simple quantitative model to further explore when user firms will find it more cost-effective to develop a solution—a new product or service—for themselves rather than hiring a manufacturer to solve the problem for them. Finally, I point out that *individual* users can sometimes be more inclined to innovate than one might expect because they sometimes value the *process* of innovating as well as the novel product or service that is created.

Users' vs. Manufacturers' Views of Innovation Opportunities

Three specific contributors to transaction costs—in addition to the “usual suspects,” such as opportunism—often have important effects on users’ decisions whether to buy a custom product or to develop it for themselves. These are (1) differences between users’ and manufacturers’ views regarding what constitutes a desirable solution, (2) differences in innovation quality signaling requirements between user and manufacturer innovators, and (3) differences in legal requirements placed on user and manufacturer innovators. The first two of these factors involve considerations of agency costs. When a user hires a manufacturer to develop a custom product, the user is a principal that has hired the custom manufacturer as to act as its agent. When the interests of the principal and the agent are not the same, agency costs will result. Recall from chapter 1 that agency costs are (1) costs incurred to monitor the agent to ensure that it follows the interests of the principal, (2) the cost incurred by the agent to commit itself not to act against the principal’s interest (the “bonding cost”), and (3) costs associated with an outcome that does not fully serve the interests of the principal (Jensen and Meckling 1976). In the specific instance of product and service development, agency considerations enter because a user’s and a manufacturer’s interests with respect to the development of a custom product often differ significantly.

Preferences Regarding Solutions

Individual products and services are components of larger user solutions. A user therefore wants a product that will make the best overall tradeoff between solution quality and price. Sometimes the best overall tradeoff will result in a willingness to pay a surprisingly large amount to get a solution component precisely right. For example, an individual user may specify tennis racket functionality that will fit her specific technique and relative strengths and will be willing to pay a great deal for exactly that racket. Deviations in racket functionality would require compensating modifications in her carefully practiced and deeply ingrained hitting technique—a much more costly overall solution from the user’s point of view. In contrast, a user will be much less concerned with precisely *how* the desired functionality is attained. For example, tennis players will typically be unconcerned

about whether a tennis racket is made from metal, carbon fiber, plastic or wood—or, for that matter, from mud—if it performs precisely as desired. And, indeed, users have quickly shifted to new types of rackets over the years as new materials promise a better fit to their particular functional requirements.

Of course, the same thing is true in the case of products for industrial users. For example, a firm with a need for a process machine may be willing to pay a great deal for one that is precisely appropriate to the characteristics of the input materials being processed, and to the skills of employees who will operate the machine. Deviations in either matter would require compensating modifications in material supply and employee training—likely to be a much more costly overall solution from the user's point of view. In contrast, the user firm will be much less concerned with precisely how the desired functionality is achieved by the process machine, and will care only that it performs precisely as specified.

Manufacturers faced with custom development requests from users make similar calculations, but theirs revolve around attempting to conserve the applicability of a low-cost (to them) solution. Manufacturers tend to specialize in and gain competitive advantage from their capabilities in one or a few specific solution types. They then seek to find as many profitable applications for those solutions types as possible. For example, a specialist in fabricating custom products from carbon fiber might find it profitable to make any kind of product—from airplane wings to tennis rackets—as long as they are made from carbon fiber. In contrast, that same manufacturer would have no competitive advantage in—and so no profit from making—any of these same products from metal or wood.

Specializations in solution types can be very narrow indeed. For example, thousands of manufacturers specialize in adhesive-based fastening solutions, while other thousands specialize in mechanical fastening solutions involving such things as metal screws and nails. Importantly, companies that produce products and solution types that have close functional equivalence from the user's point of view can look very different from the point of view of a solution supplier. For example, a manufacturer of standard or custom adhesives needs chemists on staff with an expertise in chemical formulation. It also needs chemistry labs and production equipment designed to mix specialized batches of chemicals on a small scale, and it

needs the equipment, expertise and regulatory approvals to package that kind of product in a way that is convenient to the customer and also in line with regulatory safeguards. In contrast, manufacturers specializing in standard or custom metal fastening solutions need none of these things. What they need instead are mechanical design engineers, a machine shop to build product prototypes and production tooling, specialized metal-forming production equipment such as screw machines, and so on.

Users, having an investment only in a need specification and not in a solution type, want the best functional solution to their problem, independent of solution type used. Manufacturers, in contrast, want to supply custom solutions to users that utilize their existing expertise and production capabilities. Thus, in the case of the two fastening technology alternatives just described, users will prefer whatever solution approach works best. In contrast, adhesives manufacturers will find it tremendously more attractive to create a solution involving adhesive-based fastening, and manufacturers specializing in mechanical fastening will similarly strongly prefer to offer to develop solutions involving mechanical fastening.

The difference between users' incentives to get the best functional solution to their need and specialist manufacturers' incentives to embed a specific solution type in the product to be developed are a major source of agency costs in custom product development, because there is typically an information asymmetry between user and manufacturer with respect to what will be the best solution. Manufacturers tend to know more than users about this and to have a strong incentive to provide biased information to users in order to convince them that the solution type in which they specialize is the best one to use. Such biases will be difficult for users to detect because, again, they are less expert than the suppliers in the various solution technologies that are candidates.

Theoretically, this agency cost would disappear if it were possible to fully specify a contract (Aghion and Tirole 1994; Bessen 2004). But in product development, contracting can be problematic. Information regarding characteristics of solutions and needs is inescapably incomplete at the time of contracting—users cannot fully specify what they want in advance of trying out prototype solutions, and manufacturers are not fully sure how planned solution approaches will work out before investing in customer-specific development.

Users' Expectations

When users buy a product from manufacturers, they tend to expect a package of other services to come along with the product they receive. However, when users develop a product for themselves, some of these are not demanded or can be supplied in a less formal, less expensive way by users for themselves. This set of implicit expectations can raise the cost to a user of a custom solution bought from a manufacturer relative to a home-developed solution.

Users typically expect a solution they have purchased to work correctly and reliably "right out of the box." In effect, a sharp line is drawn between product development at the manufacturer's site and routine, trouble-free usage at the purchaser's site. When the user builds a product for itself, however, both the development and the use functions are in the same organization and may explicitly be overlapped. Repeated tests and repeated repairs and improvements during early use are then more likely to be understood and tolerated as an acceptable part of the development process.

A related difference in expectations has to do with field support for a product that has been purchased rather than developed in house. In the case of a purchased custom product, users expect that manufacturers will provide replacement parts and service if needed. Responding to this expectation is costly for a custom manufacturer. It must keep a record of what it has built for each particular user, and of any special parts incorporated in that user's products so that they can be built or purchased again if needed. In contrast, if a user has developed a product for itself, it has people on site who know details of its design. These employees will be capable of rebuilding or repairing or redesigning the product *ad hoc* if and as the need arises. (Of course, if these knowledgeable employees leave the user firm while the product they designed is still in use, such informality can prove costly.)

Manufacturers also must invest in indirect quality signals that may not have an effect on actual quality, but instead are designed to assure both the specific user being served and the market in general that the product being supplied is of high quality. These represent another element of agency costs that user-innovators do not incur. When users develop an innovation for themselves, they end up intimately knowing the actual quality of the solution they have developed, and knowing why and how it is appropriate to their task. As an example, an engineer building a million-dollar process

machine for in-house use might feel it perfectly acceptable to install a precisely right and very cheap computer controller made and prominently labeled by Lego, a manufacturer of children's toys. (Lego provides computer controllers for some of its children's building kit products.) But if that same engineer saw a Lego controller in a million-dollar process machine his firm was purchasing from a specialist high-end manufacturer, he might not know enough about the design details to know that the Lego controller was precisely right for the application. In that case, the engineer and his managers might well regard the seemingly inappropriate brand name as an indirect signal of bad quality.

Manufacturers are often so concerned about a reputation for quality that they refuse to take shortcuts that a customer specifically requests and that might make sense for a particular customer, lest others get wind of what was done and take it as a negative signal about the general quality of the firm's products. For example, you may say to a maker of luxury custom cars: "I want to have a custom car of your brand in my driveway—my friends will admire it. But I only plan to drive it to the grocery store once in a while, so I only want a cheap little engine. A luxury exterior combined with cheap parts is the best solution for me in this application—just slap something together and keep the price low." The maker is likely to respond: "We understand your need, but we cannot be associated with any product of low quality. Someone else may look under the hood some day, and that would damage our reputation as a maker of fine cars. You must look elsewhere, or decide you are willing to pay the price to keep one of our fine machines idle on your driveway."

Differing Legal and Regulatory Requirements

Users that innovate do not generally face legal risk if the product they develop fails and causes costs to themselves but not to others. In contrast, manufacturers that develop and sell new products are regarded under US law as also providing an implied warranty of "fitness for the intended use." If a product does not meet this criterion, and if a different, written warranty is not in place, manufacturers can be found liable for negligence with respect to providing a defective design and failure to warn buyers (Barnes and Ulin 1984). This simple difference can cause a large difference in exposure to liability by innovators and so can drive up the costs of manufacturer-provided solutions relative to user-provided ones.

For example, a user firm that builds a novel process controller to improve its plant operations must pay its own actual costs if the self-built controller fails and ruins expensive materials being processed. On the other hand, if a controller manufacturer designed the novel controller product and sold it to customers, and a failure then occurred and could be traced back to a fault in the design, the controller manufacturer is potentially liable for actual user costs *and* punitive damages. It may also incur significant reputational losses if the unhappy user broadcasts its complaints. The logical response of a controller manufacturer to this higher risk is to charge more and/or to be much more careful with respect to running exhaustive, expensive, and lengthy tests before releasing a new product. The resulting increase in cost and delay for obtaining a manufacturer-developed product can tend to tip users toward building their own, in-house solutions.

Net Result

A net result of the foregoing considerations is that manufacturers often find that developing a custom product for only one or a few users will be unprofitable. In such cases, the transaction costs involved can make it cheaper for users with appropriate capabilities to develop the product for themselves. In larger markets, in contrast, fixed transaction costs will be spread over many customers, and the economies of scale obtainable by producing for the whole market may be substantial. In that case, it will likely be cheaper for users to buy than to innovate. As a result, manufacturers, when contacted by a user with a very specific request, will be keenly interested in how many others are likely to want this solution or elements of it. If the answer is "few," a custom manufacturer will be unlikely to accept the project.

Of course, manufacturers have an incentive to *make* markets attractive from their point of view. This can be done by deviating from precisely serving the needs of a specific custom client in order to create a solution that will be "good enough" for that client but at the same time of more interest to others. Manufacturers may do this openly by arranging meetings among custom buyers with similar needs, and then urging the group to come up with a common solution that all will find acceptable. "After all," as the representative will say, "it is clear that we cannot make a special product to suit each user, so all of you must be prepared to make really difficult compromises!" More covertly, manufacturers may simply ignore some of the

specific requests of the specific user client and make something that they expect to be a more general solution instead.

The contrasting incentives of users and manufacturers with respect to generality of need being served—and also with respect to the solution choice issue discussed earlier—can result in a very frustrating and cloudy interaction in which each party hides its best information and attempts to manipulate others to its own advantage. With respect to generality of need, sophisticated users understand custom suppliers' preference for a larger market and attempt to argue convincingly that "everyone will want precisely what I am asking you for." Manufacturers, in turn, know users have this incentive and so will generally prefer to develop custom products for which they themselves have a reasonable understanding of demand. Users are also aware of manufacturers' strong preference for only producing products that embody their existing solution expertise. To guard against the possibility that this incentive will produce biased advice, they may attempt to shop around among a number of suppliers offering different solution types and/or develop internal expertise on solution possibilities and/or attempt to write better contracts. All these attempts to induce and guard against bias involve agency costs.

An Illustrative Case

A case study by Sarah Slaughter (1993) illustrates the impact of some of the transaction costs discussed above related to users' innovate-or-buy decisions. Slaughter studied patterns of innovation in stressed-skin panels, which are used in some housing construction. The aspects of the panels studied were related to installation, and so the users of these features were home builders rather than home owners. When Slaughter contrasted users' costs of innovating versus buying, she found that it was always much cheaper for the builder to develop a solution for itself at a construction site than to ask a panel manufacturer to do so.

A stressed-skin panel can be visualized as a large 4-by-8-foot sandwich consisting of two panels made of plywood with a layer of plastic foam glued in between. The foam, about 4 inches thick, strongly bonds the two panels together and also acts as a layer of thermal insulation. In 1989, manufacturing of stressed-skin panels was a relatively concentrated industry; the four largest manufacturers collectively having a 77 percent share of the market. The user industry was much less concentrated: the four largest con-

structors of panelized housing together had only 1 percent of the market for such housing in 1989.

In housing construction, stressed-skin panels are generally attached to strong timber frames to form the outer shell of a house and to resist shear loads (such as the force of the wind). To use the panels in this way, a number of subsidiary inventions are required. For example, one must find a practical, long-lasting way to attach panels to each other and to the floors, the roof, and the frame. Also, one has to find a new way to run pipes and wires from place to place because there are no empty spaces in the walls to put them—panel interiors are filled with foam.

Stressed-skin panels were introduced into housing construction after World War II. From then till 1989, the time of Slaughter's study, 34 innovations were made in 12 functionally important areas to create a complete building system for this type of construction. Slaughter studied the history of each of these innovations and found that 82 percent had been developed by users of the stressed-skin panels—residential builders—and only 18 percent by manufacturers of stressed-skin panels. Sometimes more than one user developed and implemented different approaches to the same functional problem (table 4.1). Builders freely revealed their innovations rather than protecting them for proprietary advantage. They were passed from builder to builder by word of mouth, published in trade magazines, and diffused widely. All were replicated at building sites for years before any commercial panel manufacturer developed and sold a solution to accomplish the same function.

Histories of the user-developed improvements to stressed-skin panel construction showed that the user-innovator construction firms did not engage in planned R&D projects. Instead, each innovation was an immediate response to a problem encountered in the course of a construction project. Once a problem was encountered, the innovating builder typically developed and fabricated a solution at great speed, using skills, materials, and equipment on hand at the construction site. Builders reported that the average time from discovery of the problem to installation of the completed solution on the site was only half a day. The total cost of each innovation, including time, equipment, and materials, averaged \$153.

Example: Installing Wiring in a Stressed-Skin Panel

A builder was faced with the immediate problem of how to route wires through the foam interior of panels to wall switches located in the middle of

Table 4.1

Users would have found it much more costly to get custom solutions from manufacturers. The costs of user-developed innovations in stressed-skin panels were very low.

Function	Average user development time (days)	Average user development cost	N	Minimum cost of waiting for manufacturer to deliver
Framing of openings in panels	0.1	\$20	1	\$1,400
Structural connection between panels	0.1	30	2	\$1,400
Ventilation of panels on roof	0.1	32	2	\$28,000
Insulated connection between panels	0.1	41	3	\$2,800
Corner connection between panels	0.2	60	1	\$2,800
Installation of HVAC in panels	0.2	60	2	\$2,800
Installation of wiring in panels	0.2	79	7	\$2,800
Connection of panels to roof	0.2	80	1	\$2,800
Add insect repellency to panels	0.4	123	3	\$70,000
Connect panels to foundation	0.5	160	1	\$1,400
Connect panels to frames	1.2	377	3	\$2,800
Development of curved panels	5.0	1,500	1	\$28,000
Average for all innovations	0.5	\$153		\$12,367

N represents number of innovations developed by *users* to carry out each listed function. Source: Slaughter 1993, tables 4 and 5. Costs and times shown are averaged for all user-developed innovations in each functional category. (The six *manufacturer*-developed innovations in Slaughter’s sample are not included in this table.)

the panels. He did not want cut grooves or channels through the surfaces of the panels to these locations—that would dangerously reduce the panels’ structural strength. His inventive solution was to mount an electrically heated wire on the tip of a long pole and simply push the heated tip through the center insulation layer of the panel. As he pushed, the electrically heated tip quickly melted a channel through the foam plastic insulation from the edge of the panel to the desired spot. Wires were then pulled through this channel.

The builder-innovator reported that the total time to develop the innovation was only an hour, and that the total cost for time and materials equaled \$40. How could it cost so little and take so little time? The builder explained that using hot wires to slice sheets of plastic foam insulation into

pieces of a required length is a technique known to builders. His idea as to how to modify the slicing technique to melt channels instead came to him quickly. To test the idea, he immediately sent a worker to an electrical supply house to get some nichrome wire (a type of high-resistance wire often used as an electrical heating element), attached the wire to a tip of a pole, and tried the solution on a panel at the building site—and it worked!

This solution was described in detail in an article in a builder's magazine and was widely imitated. A panel manufacturer's eventual response (after the user solution had spread for a number of years) was to manufacture a panel with a channel for wires pre-molded into the plastic foam interior of the panel. This solution is only sometimes satisfactory. Builders often do not want to locate switch boxes at the height of the premolded channel. Also, sometimes construction workers will install some panels upside down in error, and the preformed channels will then not be continuous between one panel and the next. In such cases, the original, user-developed solution is again resorted to.

Example: Creating a Curved Panel

A builder was constructing a custom house with large, curved windows. Curved stressed-skin panels were needed to fill in the space above and below these windows, but panel manufacturers only sold flat panels at that time. The builder facing the problem could not simply buy standard flat panels and bend them into curved ones at the construction site—completed panels are rigid by design. So he bought plywood and plastic foam at a local building supply house and slowly bent each panel component separately over a curved frame quickly built at the construction site. He then bonded all three elements together with glue to create strong curved panels that would maintain their shape over time.

To determine whether users' decisions to innovate rather than buy made economic sense for them, Slaughter calculated, in a very conservative way, what it would have cost users to buy a manufacturer-developed solution embodied in a manufactured panel rather than build a solution for themselves. Her estimates included only the cost of the delay a user-builder would incur while waiting for delivery of a panel incorporating a manufacturer's solution. Delay in obtaining a solution to a problem encountered at a construction site is costly for a builder, because the schedule of deliveries,

subcontractors, and other activities must then be altered. For example, if installation of a panel is delayed, one must also reschedule the arrival of the subcontractor hired to run wires through it, the contractor hired to paint it, and so on. Slaughter estimated the cost of delay to a builder at \$280 per crew per day of delay (Means 1989). To compute delay times, she assumed that a manufacturer would always be willing to supply the special item a user requested. She also assumed that no time elapsed while the manufacturer learned about the need, contracted to do the job, designed a solution, and obtained needed regulatory approvals. She then asked panel manufacturers to estimate how long it would take them to simply construct a panel with the solution needed and deliver it to the construction site. Delay times computed in this manner ranged from 5 days for some innovations to 250 days for the longest-term one and averaged 44 days.

The conservative nature of this calculation is very clear. For example, Slaughter points out that the regulatory requirements for building components, not included, are in fact much more stringent for manufacturers than for user-builders in the field of residential construction. Manufacturers delivering products can be required to provide test data demonstrating compliance with local building codes for each locality served. Testing new products for compliance in a locality can take from a month to several years, and explicit code approval often takes several additional years. In contrast, a builder that innovates need only convince the local building inspector that what he has done meets code or performance requirements—often a much easier task (Ehrenkrantz Group 1979; Duke 1988).

Despite her very conservative method of calculation, Slaughter found the costs to users of obtaining a builder solution to be at least 100 times the actual costs of developing a solution for themselves (table 4.1). Clearly, users' decisions to innovate rather than buy made economic sense in this case.

Modeling Users' Innovate-or-Buy Decisions

In this section I summarize the core of the argument discussed in this chapter via a simple quantitative model developed with Carliss Baldwin. Our goal is to offer additional clarity by trading off the richness of the qualitative argument for simplicity.

Whether a user firm should innovate or buy is a variant of a well-known problem: where one should place an activity in a supply chain. In any real-

world case many complexities enter. In the model that follows, Baldwin and I ignore most of these and consider a simple base case focused on the impact of transaction costs on users' innovate-or-buy considerations. The model deals with manufacturing firms and user firms rather than individual users. We assume that user firms and manufacturer firms both will hire designers from the same homogeneous pool if they elect to solve a user problem. We also assume that both user firms and manufacturer firms will incur the same costs to solve a specific user problem. For example, they will have the same costs to monitor the performance of the designer employees they hire. In this way we simplify our innovate-or-buy problem to one of transaction costs only.

If there are no transaction costs (for example, no costs to write and enforce a contract), then by Coase's theorem a user will be indifferent between making or buying a solution to its problem. But in the real world there *are* transaction costs, and so a user will generally prefer to either make or buy. Which, from the point of view of minimizing overall costs of obtaining a problem solution, is the better choice under any given circumstances?

Let V_{ij} be the value of a solution to problem j for user i . Let N_j be the number of users having problem j . Let Wh_j be the cost of solving problem j , where W = hourly wage and h_j = hours required to solve it. Let P_j be the price charged by a manufacturer for a solution to problem j . Let T be fixed or "setup" transaction costs, such as writing a general contract for buyers of a solution to problem j . Let t be variable or "frictional" transaction costs, such as tailoring the general contract to a specific customer.

To explore this problem we make two assumptions. First, we assume that a user firm knows its own problems and the value of a solution to itself, V_{ij} . Second, we assume that a manufacturer knows the number of users having each problem, N_j , and the value of solutions for each problem for all users, V_{ij} .

These assumptions are in line with real-world incentives of users and manufacturers, although information stickiness generally prevents firms from getting full information. That is, users have a high incentive to know their own problems and the value to them of a solution. Manufacturers, in turn, have an incentive to invest in understanding the nature of problems faced by users in the target market, the number of users affected, and the value that the users would attach to getting a solution in order to determine the potential profitability of markets from their point of view.

We first consider the user's payoff for solving a problem for itself. A user has no transaction costs in dealing with itself, so a user's payoff for solving problem j will be $V_{ij} - Wh_j$. Therefore, a user will buy a solution from an upstream manufacturer rather than develop one for itself if and only if $P_j \leq Wh_j$.

Next we consider payoffs to a manufacturer for solving problem j . In this case, transaction costs such as those discussed in earlier sections will be encountered. With respect to transaction costs assume first that $t = 0$ but $T > 0$. Then, the manufacturer's payoff for solving problem j will be $V_{ij} - Wh_j$, which needs to be positive in order for the manufacturer to find innovation attractive:

$$N_j P_j - Wh_j - T > 0.$$

But, as we saw, $P_j \leq Wh_j$ if the user is to buy, so we may substitute Wh_j for P_j in our inequality. Thus we obtain the following inequality as a condition for the user to buy:

$$N_j(Wh_j) - Wh_j - T > 0,$$

or

$$N_j > (T / Wh_j) + 1.$$

In other words, Baldwin and I find that the absolute lower bound on N is greater than 1. This means that a single user will always prefer to solve a unique problem j for itself (except in Coase's world, where $T = 0$, and the user will be indifferent). If every problem is unique to a single user, users will never choose to call on upstream manufacturers for solutions.

Now assume that $T = 0$ but $t > 0$. Then the condition for the user to buy rather than to innovate for itself becomes

$$N_j(Wh_j - t) - Wh_j > 0,$$

or equivalently (provided $Wh_j > t$)

$$N_j > Wh_j / (Wh_j - t) > 1.$$

Again, users will not call on upstream manufacturers to solve problems unique to one user.

The findings from the simplified model, then, are the following: Problems unique to one user will always be solved efficiently by users hiring designers to work for them in house. In contrast, problems affecting more than a moderate number of users, n , which is a function of the trans-

action costs, will be efficiently solved by the manufacturer hiring designers to develop the needed new product or service and then selling that solution to all users affected by the problem. However, given sufficient levels of T and/or of t , problems affecting more than one but fewer than n users will not be solved by a manufacturer, and so there will be a market failure: Assuming an institutional framework consisting only of independent users and manufacturers, multiple users will have to solve the same problem independently.

As illustration, suppose that $t = 0.25Wh_i$ and $T = 10Wh_i$. Then, combining the two expressions and solving for n yields

$$n = (11Wh_i / 0.75Wh_i) = 14.66.$$

The condition for the user to buy the innovation rather than innovate itself becomes $N_i \geq 15$. For a number of users less than 15 but greater than 1, there will be a wasteful multiplication of user effort: several users will invest in developing the same innovation independently.

In a world that consists entirely of manufacturers and of users that do not share the innovations they develop, the type of wasteful duplicative innovation investment by users just described probably will occur often. As was discussed earlier in this chapter, and as was illustrated by Slaughter's study, substantial transaction costs might well be the norm. In addition, low numbers of users having the same need—situations where N_i is low—might also be the norm in the case of functionally novel innovations. Functionally novel innovations, as I will show later, tend to be developed by lead users, and lead users are by definition at the leading (low- N_i) edge of markets.

When the type of market failure discussed above does occur, users will have an incentive to search for institutional forms with a lower T and/or a lower t than is associated with assignment of the problem to an upstream manufacturer. One such institutional form involves interdependent innovation development among multiple users (for example, the institutional form used successfully in open source software projects that I will discuss in chapter 7). Baldwin and Clark (2003) show how this form can work to solve the problem of wasteful user innovation investments that were identified in our model. They show that, given modularity in the software's architecture, it will pay for users participating in open source software projects to generate and freely reveal some components of the needed innovation, benefiting from the fact that other users are likely to develop and reveal other

components of that innovation. At the limit, the wasteful duplication of users' innovative efforts noted above will be eliminated; each innovation component will have been developed by only one user, but will be shared by many.

Benefiting from the Innovation Process

Some individual users (not user firms) may decide to innovate for themselves rather than buy even if a traditional accounting evaluation would show that they had made a major investment in time and materials for an apparently minor reward in product functionality. The reason is that individual users may gain major rewards from the process of innovating, in addition to rewards from the product being developed. Make-or-buy evaluations typically include factors such as the time and materials that must be invested to develop a solution. These costs are then compared with the likely benefits produced by the project's "output"—the new product or service created—to determine whether the project is worth doing. This was the type of comparison made by Slaughter, for example, in assessing whether it would be better for the users to make or to buy the stressed-skin panel innovations in her sample. However, in the case of individual user-innovators, this type of assessment can provide too narrow a perspective on what actually constitutes valuable project output. Specifically, there is evidence that individuals sometimes greatly prize benefits derived from their participation in *the process* of innovation. The process, they say, can produce learning and enjoyment that is of high value to them.

In the introductory chapter, I pointed out that some recreational activities, such as solving crossword puzzles, are clearly engaged in for process rewards only: very few individuals value the end "product" of a completed puzzle. But process rewards have also been found to be important for innovators that are producing outputs that they and others do value (Hertel, Niedner, and Herrmann 2003; Lakhani and Wolf 2005). Lakhani and Wolf studied a sample of individuals ($n = 684$, response rate = 34 percent) who had written new software code and contributed it to an open source project. They asked the programmers to list their three most important reasons for doing this. Fifty-eight percent of respondents said that an important motivation for writing their code was that they had a work need (33 percent), or a non-work need (30 percent) or both (5 percent) for the code

itself. That is, they valued the project's "output" as this is traditionally viewed. However, 45 percent said that one of their top three reasons for writing code was intellectual stimulation, and 41 percent said one of their top three reasons was to improve their own programming skills (Lakhani and Wolf 2005, table 6). Elaborating on these responses, 61 percent of respondents said that their participation in the open source project was their most creative experience or was as creative as their most creative experience. Also, more than 60 percent said that "if there were one more hour in the day" they would always or often dedicate it to programming.

Csikszentmihalyi (1975, 1990, 1996) systematically studied the characteristics of tasks that individuals find intrinsically rewarding, such as rock climbing. He found that a level of challenge somewhere between boredom and fear is important, and also that the experience of "flow" gained when one is fully engaged in a task is intrinsically rewarding. Amabile (1996) proposes that intrinsic motivation is a key determining factor in creativity. She defines a creative task as one that is heuristic in nature (with no predetermined path to solution), and defines a creative outcome as a novel and appropriate (useful) response to such a task. Both conditions certainly can apply to the task of developing a product or a service.

In sum, to the extent that individual user-innovators benefit from the process of developing or modifying a product as well as from the product actually developed, they are likely to innovate even when the benefits expected from the product itself are relatively low. (Employees of a firm may wish to experience this type of intrinsic reward in their work as well, but managers and commercial constraints may give them less of an opportunity to do so. Indeed, "control over my own work" is cited by many programmers as a reason that they enjoy creating code as volunteers on open source projects more than they enjoy coding for their employers for pay.)