

1 Introduction and Overview

When I say that innovation is being democratized, I mean that users of products and services—both firms and individual consumers—are increasingly able to innovate for themselves. User-centered innovation processes offer great advantages over the manufacturer-centric innovation development systems that have been the mainstay of commerce for hundreds of years. Users that innovate can develop exactly what they want, rather than relying on manufacturers to act as their (often very imperfect) agents. Moreover, individual users do not have to develop everything they need on their own: they can benefit from innovations developed and freely shared by others.

The trend toward democratization of innovation applies to information products such as software and also to physical products. As a quick illustration of the latter, consider the development of high-performance windsurfing techniques and equipment in Hawaii by an informal user group. High-performance windsurfing involves acrobatics such as jumps and flips and turns in mid-air. Larry Stanley, a pioneer in high-performance windsurfing, described the development of a major innovation in technique and equipment to Sonali Shah:

In 1978 Jürgen Honscheid came over from West Germany for the first Hawaiian World Cup and discovered jumping, which was new to him, although Mike Horgan and I were jumping in 1974 and 1975. There was a new enthusiasm for jumping and we were all trying to outdo each other by jumping higher and higher. The problem was that . . . the riders flew off in mid-air because there was no way to keep the board with you—and as a result you hurt your feet, your legs, and the board.

Then I remembered the “Chip,” a small experimental board we had built with footstraps, and thought “it’s dumb not to use this for jumping.” That’s when I first started jumping with footstraps and discovering controlled flight. I could go so much faster than I ever thought and when you hit a wave it was like a motorcycle rider

hitting a ramp; you just flew into the air. All of a sudden not only could you fly into the air, but you could land the thing, and not only that, but you could change direction in the air!

The whole sport of high-performance windsurfing really started from that. As soon as I did it, there were about ten of us who sailed all the time together and within one or two days there were various boards out there that had footstraps of various kinds on them, and we were all going fast and jumping waves and stuff. It just kind of snowballed from there. (Shah 2000)

By 1998, more than a million people were engaged in windsurfing, and a large fraction of the boards sold incorporated the user-developed innovations for the high-performance sport.

The user-centered innovation process just illustrated is in sharp contrast to the traditional model, in which products and services are developed by manufacturers in a closed way, the manufacturers using patents, copyrights, and other protections to prevent imitators from free riding on their innovation investments. In this traditional model, a user's only role is to have needs, which manufacturers then identify and fill by designing and producing new products. The manufacturer-centric model does fit some fields and conditions. However, a growing body of empirical work shows that users are the first to develop many and perhaps most new industrial and consumer products. Further, the contribution of users is growing steadily larger as a result of continuing advances in computer and communications capabilities.

In this book I explain in detail how the emerging process of user-centric, democratized innovation works. I also explain how innovation by users provides a very necessary complement to and feedstock for manufacturer innovation.

The ongoing shift of innovation to users has some very attractive qualities. It is becoming progressively easier for many users to get precisely what they want by designing it for themselves. And innovation by users appears to increase social welfare. At the same time, the ongoing shift of product-development activities from manufacturers to users is painful and difficult for many manufacturers. Open, distributed innovation is "attacking" a major structure of the social division of labor. Many firms and industries must make fundamental changes to long-held business models in order to adapt. Further, governmental policy and legislation sometimes preferentially supports innovation by manufacturers. Considerations of social welfare suggest that this must change. The workings of the intellec-

tual property system are of special concern. But despite the difficulties, a democratized and user-centric system of innovation appears well worth striving for.

Users, as the term will be used in this book, are firms or individual consumers that expect to benefit from *using* a product or a service. In contrast, manufacturers expect to benefit from *selling* a product or a service. A firm or an individual can have different relationships to different products or innovations. For example, Boeing is a manufacturer of airplanes, but it is also a user of machine tools. If we were examining innovations developed by Boeing for the airplanes it sells, we would consider Boeing a manufacturer-innovator in those cases. But if we were considering innovations in metal-forming machinery developed by Boeing for in-house use in building airplanes, we would categorize those as user-developed innovations and would categorize Boeing as a user-innovator in those cases.

Innovation user and innovation manufacturer are the two general “functional” relationships between innovator and innovation. Users are unique in that they alone benefit *directly* from innovations. All others (here lumped under the term “manufacturers”) must sell innovation-related products or services to users, indirectly or directly, in order to profit from innovations. Thus, in order to profit, inventors must sell or license knowledge related to innovations, and manufacturers must sell products or services incorporating innovations. Similarly, suppliers of innovation-related materials or services—unless they have direct use for the innovations—must sell the materials or services in order to profit from the innovations.

The user and manufacturer categorization of relationships between innovator and innovation can be extended to specific functions, attributes, or features of products and services. When this is done, it may turn out that different parties are associated with different attributes of a particular product or service. For example, householders are the users of the switching attribute of a household electric light switch—they use it to turn lights on and off. However, switches also have other attributes, such as “easy wiring” qualities, that may be used only by the electricians who install them. Therefore, if an electrician were to develop an improvement to the installation attributes of a switch, it would be considered a user-developed innovation.

A brief overview of the contents of the book follows.

Development of Products by Lead Users (Chapter 2)

Empirical studies show that many users—from 10 percent to nearly 40 percent—engage in developing or modifying products. About half of these studies do not determine representative innovation frequencies; they were designed for other purposes. Nonetheless, when taken together, the findings make it very clear that users are doing a *lot* of product modification and product development in many fields.

Studies of innovating users (both individuals and firms) show them to have the characteristics of “lead users.” That is, they are ahead of the majority of users in their populations with respect to an important market trend, and they expect to gain relatively high benefits from a solution to the needs they have encountered there. The correlations found between innovation by users and lead user status are highly significant, and the effects are very large.

Since lead users are at the leading edge of the market with respect to important market trends, one can guess that many of the novel products they develop for their own use will appeal to other users too and so might provide the basis for products manufacturers would wish to commercialize. This turns out to be the case. A number of studies have shown that many of the innovations reported by lead users are judged to be commercially attractive and/or have actually been commercialized by manufacturers.

Research provides a firm grounding for these empirical findings. The two defining characteristics of lead users and the likelihood that they will develop new or modified products have been found to be highly correlated (Morrison et al. 2004). In addition, it has been found that the higher the intensity of lead user characteristics displayed by an innovator, the greater the commercial attractiveness of the innovation that the lead user develops (Franke and von Hippel 2003a). In figure 1.1, the increased concentration of innovations toward the right indicates that the likelihood of innovating is higher for users having higher lead user index values. The rise in average innovation attractiveness as one moves from left to right indicates that innovations developed by lead users tend to be more commercially attractive. (Innovation attractiveness is the sum of the novelty of the innovation and the expected future generality of market demand.)




Figure 1.1

User-innovators with stronger “lead user” characteristics develop innovations having higher appeal in the general marketplace. Estimated OLS function: $Y = 2.06 + 0.57x$, where Y represents attractiveness of innovation and x represents lead-user-ness of respondent. Adjusted $R^2 = 0.281$; $p = 0.002$; $n = 30$. Source of data: Franke and von Hippel 2003.

Why Many Users Want Custom Products (Chapter 3)

Why do so many users develop or modify products for their own use? Users may innovate if and as they want something that is not available on the market and are able and willing to pay for its development. It is likely that many users do not find what they want on the market. Meta-analysis of market-segmentation studies suggests that users’ needs for products are highly heterogeneous in many fields (Franke and Reisinger 2003).

Mass manufacturers tend to follow a strategy of developing products that are designed to meet the needs of a large market segment well enough to induce purchase from and capture significant profits from a large number of customers. When users’ needs are heterogeneous, this strategy of “a few sizes fit all” will leave many users somewhat dissatisfied with the commercial products on offer and probably will leave some users seriously dissatisfied. In a study of a sample of users of the security features of Apache web server software, Franke and von Hippel (2003b) found that users had a very high heterogeneity of need, and that many had a high willingness to pay to

get precisely what they wanted. Nineteen percent of the users sampled actually innovated to tailor Apache more closely to their needs. Those who did were found to be significantly more satisfied.

Users' Innovate-or-Buy Decisions (Chapter 4)

Even if many users want “exactly right products” and are willing and able to pay for their development, why do users often do this for themselves rather than hire a custom manufacturer to develop a special just-right product for them? After all, custom manufacturers specialize in developing products for one or a few users. Since these firms are specialists, it is possible that they could design and build custom products for individual users or user firms faster, better, or cheaper than users could do this for themselves. Despite this possibility, several factors can drive users to innovate rather than buy. Both in the case of user firms and in the case of individual user-innovators, agency costs play a major role. In the case of individual user-innovators, enjoyment of the innovation process can also be important.

With respect to agency costs, consider that when a user develops its own custom product that user can be trusted to act in its own best interests. When a user hires a manufacturer to develop a custom product, the situation is more complex. The user is then a principal that has hired the custom manufacturer to act as its agent. If the interests of the principal and the agent are not the same, there will be agency costs. In general terms, agency costs are (1) costs incurred to monitor the agent to ensure that it (or he or she) 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, a major divergence of interests between user and custom manufacturer does exist: the user wants to get precisely what it needs, to the extent that it can afford to do so. In contrast, the custom manufacturer wants to lower its development costs by incorporating solution elements it already has or that it predicts others will want in the future—even if by doing so it does not serve its present client's needs as well as it could.

A user wants to preserve its need specification because that specification is chosen to make *that user's* overall solution quality as high as possible at

the desired price. For example, an individual user may specify a mountain-climbing boot that will precisely fit his unique climbing technique and allow him to climb Everest more easily. Any deviations in boot design will require compensating modifications in the climber's carefully practiced and deeply ingrained climbing technique—a much more costly solution from the user's point of view. A custom boot manufacturer, in contrast, will have a strong incentive to incorporate the materials and processes it has in stock and expects to use in future even if this produces a boot that is not precisely right for the present customer. For example, the manufacturer will not want to learn a new way to bond boot components together even if that would produce the best custom result for one client. The net result is that when one or a few users want something special they will often get the best result by innovating for themselves.

A small model of the innovate-or-buy decision follows. This model shows in a quantitative way that user firms with unique needs will always be better off developing new products for themselves. It also shows that development by manufacturers can be the most economical option when n or more user firms want the same thing. However, when the number of user firms wanting the same thing falls between 1 and n , manufacturers may not find it profitable to develop a new product for just a few users. In that case, more than one user may invest in developing the same thing independently, owing to market failure. This results in a waste of resources from the point of view of social welfare. The problem can be addressed by new institutional forms, such as the user innovation communities that will be studied later in this book.

Chapter 4 concludes by pointing out that an additional incentive can drive individual user-innovators to innovate rather than buy: they may value the *process* of innovating because of the enjoyment or learning that it brings them. It might seem strange that user-innovators can enjoy product development enough to want to do it themselves—after all, manufacturers pay their product developers to do such work! On the other hand, it is also clear that enjoyment of problem solving is a motivator for many individual problem solvers in at least some fields. Consider for example the millions of crossword-puzzle aficionados. Clearly, for these individuals enjoyment of the problem-solving process rather than the solution is the goal. One can easily test this by attempting to offer a puzzle solver a completed puzzle—the very output he or she is working so hard to create. One will very likely

be rejected with the rebuke that one should not spoil the fun! Pleasure as a motivator can apply to the development of commercially useful innovations as well. Studies of the motivations of volunteer contributors of code to widely used software products have shown that these individuals too are often strongly motivated to innovate by the joy and learning they find in this work (Hertel et al. 2003; Lakhani and Wolf 2005).

Users' Low-Cost Innovation Niches (Chapter 5)

An exploration of the basic processes of product and service development show that users and manufacturers tend to develop different *types* of innovations. This is due in part to information asymmetries: users and manufacturers tend to know different things. Product developers need two types of information in order to succeed at their work: need and context-of-use information (generated by users) and generic solution information (often initially generated by manufacturers specializing in a particular type of solution). Bringing these two types of information together is not easy. Both need information and solution information are often very “sticky”—that is, costly to move from the site where the information was generated to other sites. As a result, users generally have a more accurate and more detailed model of their needs than manufacturers have, while manufacturers have a better model of the solution approach in which they specialize than the user has.

When information is sticky, innovators tend to rely largely on information they already have in stock. One consequence of the information asymmetry between users and manufacturers is that users tend to develop innovations that are functionally novel, requiring a great deal of user-need information and use-context information for their development. In contrast, manufacturers tend to develop innovations that are improvements on well-known needs and that require a rich understanding of solution information for their development. For example, firms that use inventory-management systems, such as retailers, tend to be the developers of new approaches to inventory management. In contrast, manufacturers of inventory-management systems and equipment tend to develop improvements to the equipment used to implement these user-devised approaches (Ogawa 1998).

If we extend the information-asymmetry argument one step further, we see that information stickiness implies that information on hand will also differ among *individual* users and manufacturers. The information assets of

some particular user (or some particular manufacturer) will be closest to what is required to develop a particular innovation, and so the cost of developing that innovation will be relatively low for that user or manufacturer. The net result is that user innovation activities will be *distributed* across many users according to their information endowments. With respect to innovation, one user is by no means a perfect substitute for another.

Why Users Often Freely Reveal Their Innovations (Chapter 6)

The social efficiency of a system in which individual innovations are developed by individual users is increased if users somehow diffuse what they have developed to others. Manufacturer-innovators *partially* achieve this when they sell a product or a service on the open market (partially because they diffuse the product incorporating the innovation, but often not all the information that others would need to fully understand and replicate it). If user-innovators do not somehow also diffuse what they have done, multiple users with very similar needs will have to independently develop very similar innovations—a poor use of resources from the viewpoint of social welfare. Empirical research shows that users often do achieve widespread diffusion by an unexpected means: they often “freely reveal” what they have developed. When we say that an innovator freely reveals information about a product or service it has developed, we mean that all intellectual property rights to that information are voluntarily given up by the innovator, and all interested parties are given access to it—the information becomes a public good.

The empirical finding that users often freely reveal their innovations has been a major surprise to innovation researchers. On the face of it, if a user-innovator’s proprietary information has value to others, one would think that the user would strive to prevent free diffusion rather than help others to free ride on what it has developed at private cost. Nonetheless, it is now very clear that individual users and user firms—and sometimes manufacturers—often freely reveal detailed information about their innovations.

The practices visible in “open source” software development were important in bringing this phenomenon to general awareness. In these projects it was clear *policy* that project contributors would routinely and systematically freely reveal code they had developed at private expense (Raymond 1999). However, free revealing of product innovations has a history that began

long before the advent of open source software. Allen, in his 1983 study of the eighteenth-century iron industry, was probably the first to consider the phenomenon systematically. Later, Nuvolari (2004) discussed free revealing in the early history of mine pumping engines. Contemporary free revealing by users has been documented by von Hippel and Finkelstein (1979) for medical equipment, by Lim (2000) for semiconductor process equipment, by Morrison, Roberts, and von Hippel (2000) for library information systems, and by Franke and Shah (2003) for sporting equipment. Henkel (2003) has documented free revealing among manufacturers in the case of embedded Linux software.

Innovators often freely reveal because it is often the best or the only practical option available to them. Hiding an innovation as a trade secret is unlikely to be successful for long: too many generally know similar things, and some holders of the “secret” information stand to lose little or nothing by freely revealing what they know. Studies find that innovators in many fields view patents as having only limited value. Copyright protection and copyright licensing are applicable only to “writings,” such as books, graphic images, and computer software.

Active efforts by innovators to freely reveal—as opposed to sullen acceptance—are explicable because free revealing can provide innovators with significant private benefits as well as losses or risks of loss. Users who freely reveal what they have done often find that others then improve or suggest improvements to the innovation, to mutual benefit (Raymond 1999). Freely revealing users also may benefit from enhancement of reputation, from positive network effects due to increased diffusion of their innovation, and from other factors. Being the first to freely reveal a particular innovation can also enhance the benefits received, and so there can actually be a rush to reveal, much as scientists rush to publish in order to gain the benefits associated with being the first to have made a particular advancement.

Innovation Communities (Chapter 7)

Innovation by users tends to be widely distributed rather than concentrated among just a very few very innovative users. As a result, it is important for user-innovators to find ways to combine and leverage their efforts. Users achieve this by engaging in many forms of cooperation. Direct, informal user-to-user cooperation (assisting others to innovate, answering questions,

and so on) is common. Organized cooperation is also common, with users joining together in networks and communities that provide useful structures and tools for their interactions and for the distribution of innovations. Innovation communities can increase the speed and effectiveness with which users and also manufacturers can develop and test and diffuse their innovations. They also can greatly increase the ease with which innovators can build larger systems from interlinkable modules created by community participants.

Free and open source software projects are a relatively well-developed and very successful form of Internet-based innovation community. However, innovation communities are by no means restricted to software or even to information products, and they can play a major role in the development of physical products. Franke and Shah (2003) have documented the value that user innovation communities can provide to user-innovators developing physical products in the field of sporting equipment. The analogy to open source innovation communities is clear.

The collective or community effort to provide a public good—which is what freely revealed innovations are—has traditionally been explored in the literature on “collective action.” However, behaviors seen in extant innovation communities fail to correspond to that literature at major points. In essence, innovation communities appear to be more robust with respect to recruiting and rewarding members than the literature would predict. Georg von Krogh and I attribute this to innovation contributors’ obtaining some private rewards that are not shared equally by free riders (those who take without contributing). For example, a product that a user-innovator develops and freely reveals might be perfectly suited to that user-innovator’s requirements but less well suited to the requirements of free riders. Innovation communities thus illustrate a “private-collective” model of innovation incentive (von Hippel and von Krogh 2003).

Adapting Policy to User Innovation (Chapter 8)

Is innovation by users a “good thing?” Welfare economists answer such a question by studying how a phenomenon or a change affects social welfare. Henkel and von Hippel (2005) explored the social welfare implications of user innovation. They found that, relative to a world in which only manufacturers innovate, social welfare is very probably increased by the presence

of innovations freely revealed by users. This finding implies that policy making should support user innovation, or at least should ensure that legislation and regulations do not favor manufacturers at the expense of user-innovators.

The transitions required of policy making to achieve neutrality with respect to user innovation vs. manufacturer innovation are significant. Consider the impact on open and distributed innovation of past and current policy decisions. Research done in the past 30 years has convinced many academics that intellectual property law is sometimes or often not having its intended effect. Intellectual property law was intended to increase the amount of innovation investment. Instead, it now appears that there are economies of scope in both patenting and copyright that allow firms to use these forms of intellectual property law in ways that are directly opposed to the intent of policy makers and to the public welfare. Major firms can invest to develop large portfolios of patents. They can then use these to create “patent thickets”—dense networks of patent claims that give them plausible grounds for threatening to sue across a wide range of intellectual property. They may do this to prevent others from introducing a superior innovation and/or to demand licenses from weaker competitors on favorable terms (Shapiro 2001). Movie, publishing, and software firms can use large collections of copyrighted work to a similar purpose (Benkler 2002). In view of the distributed nature of innovation by users, with each tending to create a relatively small amount of intellectual property, users are likely to be disadvantaged by such strategies.

It is also important to note that users (and manufacturers) tend to build prototypes of their innovations economically by modifying products already available on the market to serve a new purpose. Laws such as the (US) Digital Millennium Copyright Act, intended to prevent consumers from illegally copying protected works, also can have the unintended side effect of preventing users from modifying products that they purchase (Varian 2002). Both fairness and social welfare considerations suggest that innovation-related policies should be made neutral with respect to the sources of innovation.

It may be that current impediments to user innovation will be solved by legislation or by policy making. However, beneficiaries of existing law and policy will predictably resist change. Fortunately, a way to get around some of these problems is in the hands of innovators themselves. Suppose many

innovators in a particular field decide to freely reveal what they have developed, as they often have reason to do. In that case, users can collectively create an information commons (a collection of information freely available to all) containing substitutes for some or a great deal of information now held as private intellectual property. Then user-innovators can work around the strictures of intellectual property law by simply using these freely revealed substitutes (Lessig 2001). This is essentially what is happening in the field of software. For many problems, user-innovators in that field now have a choice between proprietary, closed software provided by Microsoft and other firms and open source software that they can legally download from the Internet and legally modify to serve their own specific needs.

Policy making that levels the playing field between users and manufacturers will force more rapid change onto manufacturers but will by no means destroy them. Experience in fields where open and distributed innovation processes are far advanced show how manufacturers can and do adapt. Some, for example, learn to supply proprietary platform products that offer user-innovators a framework upon which to develop and use their improvements.

Democratizing Innovation (Chapter 9)

Users' ability to innovate is improving *radically* and *rapidly* as a result of the steadily improving quality of computer software and hardware, improved access to easy-to-use tools and components for innovation, and access to a steadily richer innovation commons. Today, user firms and even individual hobbyists have access to sophisticated programming tools for software and sophisticated CAD design tools for hardware and electronics. These information-based tools can be run on a personal computer, and they are rapidly coming down in price. As a consequence, innovation by users will continue to grow even if the degree of heterogeneity of need and willingness to invest in obtaining a precisely right product remains constant.

Equivalents of the innovation resources described above have long been available within corporations to a few. Senior designers at firms have long been supplied with engineers and designers under their direct control, and with the resources needed to quickly construct and test prototype designs. The same is true in other fields, including automotive design and clothing

design: just think of the staffs of engineers and modelmakers supplied so that top auto designers can quickly realize and test their designs.

But if, as we have seen, the information needed to innovate in important ways is widely distributed, the traditional pattern of concentrating innovation-support resources on a few individuals is hugely inefficient. High-cost resources for innovation support cannot efficiently be allocated to “the right people with the right information:” it is very difficult to know who these people may be before they develop an innovation that turns out to have general value. When the cost of high-quality resources for design and prototyping becomes very low (the trend we have described), these resources can be diffused very widely, and the allocation problem diminishes in significance. The net result is and will be to democratize the opportunity to create.

On a level playing field, users will be an increasingly important source of innovation and will increasingly substitute for or complement manufacturers’ innovation-related activities. In the case of information products, users have the possibility of largely or completely doing without the services of manufacturers. Open source software projects are object lessons that teach us that users can create, produce, diffuse, provide user field support for, update, and use complex products by and for themselves in the context of user innovation communities. In physical product fields, product development by users can evolve to the point of largely or totally supplanting product development—but not product manufacturing—by manufacturers. (The economies of scale associated with manufacturing and distributing physical products give manufacturers an advantage over “do-it-yourself” users in those activities.)

The evolving pattern of the locus of product development in kitesurfing illustrates how users can displace manufacturers from the role of product developer. In that industry, the collective product-design and testing work of a user innovation community has clearly become superior in both quality and quantity relative to the levels of in-house development effort that manufacturers of kitesurfing equipment can justify. Accordingly, manufacturers of such equipment are increasingly shifting away from product design and focusing on producing product designs first developed and tested by user innovation communities.

How can or should manufacturers adapt to users’ encroachment on elements of their traditional business activities? There are three general possi-

bilities: (1) Produce user-developed innovations for general commercial sale and/or offer custom manufacturing to specific users. (2) Sell kits of product-design tools and/or “product platforms” to ease users’ innovation-related tasks. (3) Sell products or services that are complementary to user-developed innovations. Firms in fields where users are already very active in product design are experimenting with all these possibilities.

Application: Searching for Lead User Innovations (Chapter 10)

Manufacturers design their innovation processes around the way they think the process works. The vast majority of manufacturers still think that product development and service development are always done by manufacturers, and that their job is always to find a need and fill it rather than to sometimes find and commercialize an innovation that lead users have already developed. Accordingly, manufacturers have set up market-research departments to explore the needs of users in the target market, product-development groups to think up suitable products to address those needs, and so forth. The needs and prototype solutions of lead users—if encountered at all—are typically rejected as outliers of no interest. Indeed, when lead users’ innovations do enter a firm’s product line—and they have been shown to be the actual source of many major innovations for many firms—they typically arrive with a lag and by an unconventional and unsystematic route. For example, a manufacturer may “discover” a lead user innovation only when the innovating user firm contacts the manufacturer with a proposal to produce its design in volume to supply its own in-house needs. Or sales or service people employed by a manufacturer may spot a promising prototype during a visit to a customer’s site.

Modification of firms’ innovation processes to *systematically* search for and further develop innovations created by lead users can provide manufacturers with a better interface to the innovation process as it actually works, and so provide better performance. A natural experiment conducted at 3M illustrates this possibility. Annual sales of lead user product ideas generated by the average lead user project at 3M were conservatively forecast by management to be more than 8 times the sales forecast for new products developed in the traditional manner—\$146 million versus \$18 million per year. In addition, lead user projects were found to generate ideas for new product lines, while traditional market-research methods were found to

produce ideas for incremental improvements to existing product lines. As a consequence, 3M divisions funding lead user project ideas experienced their highest rate of major product line generation in the past 50 years (Lilien et al. 2002).

Application: Toolkits for User Innovation and Custom Design (Chapter 11)

Firms that understand the distributed innovation process and users' roles in it can *change* factors affecting lead user innovation and so affect its rate and direction in ways they value. Toolkits for user innovation custom design offer one way of doing this. This approach involves partitioning product-development and service-development projects into solution-information-intensive subtasks and need-information-intensive subtasks. Need-intensive subtasks are then assigned to users along with a kit of tools that enable them to effectively execute the tasks assigned to them. The resulting co-location of sticky information and problem-solving activity makes innovation within the solution space offered by a particular toolkit cheaper for users. It accordingly attracts them to the toolkit and so influences what they develop and how they develop it. The custom semiconductor industry was an early adopter of toolkits. In 2003, more than \$15 billion worth of semiconductors were produced that had been designed using this approach.

Manufacturers that adopt the toolkit approach to supporting and channeling user innovation typically face major changes in their business models, and important changes in industry structure may also follow. For example, as a result of the introduction of toolkits to the field of semiconductor manufacture, custom semiconductor manufacturers—formerly providers of both design and manufacturing services to customers—lost much of the work of custom product design to customers. Many of these manufacturers then became specialist silicon foundries, supplying production services primarily. Manufacturers may or may not wish to make such changes. However, experience in fields where toolkits have been deployed shows that customers tend to prefer designing their own custom products with the aid of a toolkit over traditional manufacturer-centric development practices. As a consequence, the only real choice for manufacturers in a field appropriate to the deployment of toolkits may be whether to lead or to follow in the transition to toolkits.

Linking User Innovation to Other Phenomena and Fields (Chapter 12)

In chapter 12 I discuss links between user innovation and some related phenomena and literatures. With respect to phenomena, I point out the relationship of user innovation to *information* communities, of which user innovation communities are a subset. One open information community is the online encyclopedia Wikipedia (www.wikipedia.org). Other such communities include the many specialized Internet sites where individuals with both common and rare medical conditions can find one another and can find specialists in those conditions. Many of the advantages associated with user innovation communities also apply to open information networks and communities. Analyses appropriate to information communities follow the same overall pattern as the analyses provided in this book for innovation communities. However, they are also simpler, because in open information communities there may be little or no proprietary information being transacted and thus little or no risk of related losses for participants.

Next I discuss links between user-centric innovation phenomena and the literature on the economics of knowledge that have been forged by Foray (2004) and Weber (2004). I also discuss how Porter's 1991 work on the competitive advantage of nations can be extended to incorporate findings on nations' lead users as product developers. Finally, I point out how findings explained in this book link to and complement research on the Social Construction of Technology (Pinch and Bijker 1987).

I conclude this introductory chapter by reemphasizing that user innovation, free revealing, and user innovation communities will flourish under many but not all conditions. What we know about manufacturer-centered innovation is still valid; however, lead-user-centered innovation patterns are increasingly important, and they present major new opportunities and challenges for us all.

