

Profiting from voluntary information spillovers:
How users benefit
by freely revealing their innovations

Dietmar Harhoff^{1,2}, Joachim Henkel^{1,3} and Eric von Hippel³

¹ University of Munich

² Centre for Economic Policy Research (CEPR), London

³ M.I.T. Sloan School of Management

May 2002

Abstract

Empirical studies of innovation have found that end users frequently develop important product and process innovations. Defying conventional wisdom on the negative effects of uncompensated spillovers, innovative users also often openly reveal their innovations to competing users and to manufacturers. Rival users are thus in a position to reproduce the innovation in-house and benefit from using it, and manufacturers are in a position to refine the innovation and sell it to all users, including competitors of the user revealing its innovation. In this paper we explore the incentives that users might have to freely reveal their proprietary innovations. We then develop a game-theoretic model to explore the effect of these incentives on users' decisions to reveal or hide their proprietary information. We find that, under realistic parameter constellations, free revealing pays. We conclude by discussing some implications of our findings.

Keywords: innovation, spillovers, diffusion, lead users

1 Introduction

A large body of research has shown that users of products and processes are the developers of many important innovations that are later produced and sold by manufacturers (e.g., Enos 1962, Knight 1963, Freeman 1968, Shaw 1985, von Hippel 1988). The analysis of these innovation processes has also often shown a puzzling phenomenon: Innovating users often do not sell or license their innovations to manufacturers. Instead, they *freely reveal* details of their innovations to other users and to manufacturers (von Hippel and Finkelstein 1979, Allen 1983, Morrison et al. 2000, Lim, 2000, Lakhani and von Hippel 2000). Free revealing is also a central feature of “open source” software development projects. Contributors to such projects freely reveal the novel software code they have developed to fellow innovators and to free riders on equal terms (e.g., Raymond 1999, Lerner and Tirole 2002).

When we say that an innovator “freely reveals” proprietary information, we mean that all existing and potential intellectual property rights to that information are voluntarily given up by that innovator and all interested parties are given access to it – the information becomes a public good. Thus, we define the free revealing of information by a possessor as the granting of access to all interested agents without imposition of any direct payment. For example, placement of non-patented information in a publicly-accessible site such as a journal or public website would be free revealing as we define it.

Free revealing as so defined does not mean that recipients necessarily acquire and utilize the revealed information at no cost to themselves. Recipients may, for example, have to pay for a journal subscription or an Internet connection or a field trip to acquire the information being freely revealed. Also, some may have to obtain complementary information or other assets in order to fully understand that information or put it to use. However, if the information possessor does not profit from any such expenditures made by information adopters, the information itself is still freely revealed, according to our definition.

Conversely, note that innovators may sometimes choose to subsidize the acquisition, evaluation and use of their freely-revealed information by others. For example, a firm may invest in extensive and expensive lobbying to get others to adopt a technical standard it has developed. Similarly, writers of computer code that they freely reveal may work hard to document their code in a way that is very easy for potential adopters to understand. Such subsidization efforts do not affect the status of the information itself as being freely revealed according to our definition.

To economists, free revealing is surprising, because it violates a central tenant of the economic theory of innovation. In this classical view, appropriating returns to inno-

vation requires agents to keep the knowledge underlying an innovation secret or to protect it by patents (or other means). After all, non-compensated spillovers of innovation-related information should represent a loss that innovators would seek to avoid if at all possible, even at some cost. Why then do we observe that some innovation-related information is revealed freely?

In this paper we explore this question by first describing incentives which might induce user-innovators to freely reveal their innovations. We summarize explanations that have been proposed in the literature and we add some of our own ideas as well. Next, we link these incentives to qualitative evidence regarding the free-revealing of innovations by users in a number of fields (section 2). Then we use a game-theoretic model to explore the effect of the incentives and derive conditions under which it might pay users to freely reveal their innovations to other users and/or to manufacturers. We conclude that, under very plausible parameter constellations, it does indeed pay users to freely reveal their innovations to other users – even to direct rivals (section 3). Finally, we discuss the implications of free revealing of innovations, suggesting that it is a crucial element in an emergent innovation process characterized by a distribution of innovation-related activities between users and manufacturers. We also discuss possible extensions of our work in the final section of the paper (section 4).

2 Incentives and qualitative evidence for free revealing

2.1 Review of the literature

The specific context in which we explore free revealing in this paper involves users that have developed an innovation for in-house use. In principle, such users have a choice among three options: they may keep the innovation secret, may license it, or they may choose to freely reveal it. In this paper we contrast the costs and benefits of only two choices, i.e., of secrecy versus free revealing. We do this because a number of empirical studies conducted by several authors over a span of many years have found the patent grant may be crucial in excluding rivals from using the patented technology, but not very useful for capturing royalty income in almost all industries. The few exceptions to this rule typically found are in the fields of pharmaceuticals, chemicals and chemical processes. In almost all industrial fields, intellectual property subject to licensing is protected by patents. The notable exception is software, where material to be licensed is often protected by copyright.

Literature on returns from licensing patents includes a study by Taylor and Silberston (1973). They examined the impact of British and foreign patents in a very rich study of 44 British and multinational firms selected from five broad “classes” of industrial activity. They found that these firms gained relatively little from licensing, when

benefits were computed as licensing fees and/or other considerations received minus patenting and licensing costs incurred by the innovating firm. Wilson (1975) studied data on royalty payments submitted by some U.S. corporations to the U.S. Securities and Exchange Commission in 1971 on Form 10K. He too, found corporate returns from licensing to be generally low.

The low returns from licensing patented knowledge found by Taylor and Silberston and by Wilson could be caused by weakness in protection afforded by patents or by patent holders' disinclination to license. A second type of finding, however, suggests that protection afforded by the patent system is itself generally weak, and that innovators in most fields probably could not expect to benefit from licensing their patented knowledge even if they chose to do so. A study by Scherer (1959) found only eight of thirty seven respondents ("executives responsible for technical change") reporting that patents were "very important" to their companies. Fourteen reported that patents had "some importance", and fifteen said they were "not very important" to their firms (Scherer 1959, 117). This result is especially interesting because Scherer selected his sample only from the firms which presumably value patents most highly – those which hold a large number of them. A similar finding was reported by Taylor and Silberston (1973). They report that 24 of the 32 responding firms said that 5 percent or less of recent R&D expenditures would not have been undertaken if patent protection had not been available (Taylor and Silberston 1973, 30). Levin et al., in a survey of 650 R&D executives in 130 different industries found that all except respondents from the chemical and pharmaceutical industries judged patents to be "relatively ineffective". Similar findings are reported by Mansfield (1986, 1986) and reaffirmed by Cohen et al. (2000).

Finally, we note that patents may cost a significant amount in both time and expenditure to obtain. Therefore, the protection they can afford to intellectual property developers in the few industries where they are effective will be useful only for larger innovations that are expected to retain their value for a number of years. (The average patent grant at the European Patent Office involves patent office fees of €4,300 in 2000. Professional representation before the EPO amounts to €5,500 on average and translation into the languages of eight contracting states requires €11,500. Renewal fees for a patent maintained for ten years amount to roughly €8,500.¹ For patent applications via the PCT procedure, much higher costs apply.

Given that licensing is seldom a practical option, innovating users are typically left with a choice between two alternatives: should they freely reveal their innovation to

1 See "Cost of an average European patent", http://www.european-patent-office.org/epo/new/kosten_e.pdf (Jan. 14, 2002).

a manufacturer and all others – including competing users - or should they (try to) keep it secret while profiting from in-house use?

Interestingly, the ability to keep intellectual property secret is not a general nor necessarily a long-term possibility. Much intellectual property cannot be kept a secret because, as in the case of products placed into the open market, it must be revealed in order to be sold. It is true that user-innovators are in a position to benefit from their process innovations while keeping them secret behind their factory walls. All innovators are also in principle able to keep their innovations secret while developing them and before putting them on the market. Even within these categories, however, Mansfield (1985) has found that the period during which intellectual property can be kept secret is quite limited. He studied a sample of 100 American firms and found that “...information concerning development decisions is generally in the hands of rivals within about 12 to 18 months, on the average, and information concerning the detailed nature and operation of a new product or process generally leaks out within about a year.” This suggests that the true choice facing most user-innovators is the choice between secrecy or selective or free revealing for a period of about a year, and only between voluntary or involuntary free revealing in the longer run.

To our knowledge, the literature on what we are calling free revealing as an economically-rational behavior among profit-seeking firms begins with Allen (1983). Allen described a phenomenon he calls “collective invention” that he observed in archival materials related to the nineteenth-century English iron industry. Two important attributes of iron furnaces used by that industry were subject to steady improvement during 1850-1875 in England’s Cleveland district: an increase in the height of furnace chimneys, and an increase in the temperature of the “blast” air pumped into an iron furnace during operation. Both types of technical change resulted in a significant and progressive improvement in the energy efficiency of iron production. Allen found that in the technical writings of the time, at least some innovators publicly revealed data on their furnace design and performance in meetings of professional societies and in published material.

Allen suggested three reasons that might provide an economic rationalization for this pattern of behavior by profit-seeking firms. First, he proposed that reputation gained for the firm or firm managers might offset a reduction in profits for the firm caused by free revealing. Second, he reasoned that so many people might know the information that it “...would have been costly to keep it secret. In the case of blast furnaces and steelworks, the construction would have been done by contractors who would know the design. The designs themselves were often created by consulting engineers who shifted from firm to firm (Allen p.17).” Third, he proposed that the behavior might have effects that actually increase firm profits. This may be so because the revealed innovation is to some degree specific to assets owned by the innovator (see also

Hirschleifer 1971) or to certain features of his production process (which may be shared by some other firms, in particular those in the same region, but not by the whole industry). Revealing may also increase the innovator's profit by enlarging the overall market for the product under consideration.

De Fraja (1996) demonstrates that two firms engaged in a patent race may disclose information voluntarily if the payoff structure does not allocate the full gain to the winner. In this case, being second quickly may be preferable to being first relatively late; hence, firms allow for voluntary spillovers in order to accelerate the race.

Mishina (1989) studied free revealing in the lithographic equipment industry. He showed that firms may sometimes prefer to reveal their innovation to suppliers even if doing so entails a loss of comparative advantage. This loss may be dominated by increases in productivity coming from cost reductions triggered by the new equipment. To illustrate this source of incentive, consider that, by freely revealing information regarding an innovative product or process, a user makes it possible for manufacturers to learn about that innovation, to improve upon it and also possibly to offer it at a lower price than the cost incurred by users seeking to produce it on a smaller scale in-house. When the improved version is offered for sale to the general market, the original user-innovator (and other users as well) is able to acquire it and gain from in-house use of the improvements. For example, consider that manufacturers often convert user-developed innovations ("home-builts") into a much more robust and reliable form when preparing them for sale on the commercial market. Also, manufacturers offer related services, such as field maintenance and repair programs, that innovating users must otherwise provide for themselves.

A number of researchers have explored the incentives that might induce open source software developers to freely reveal the source code they have produced. One incentive that has been proposed as relevant in this context is the reputational incentive that was posited with respect to collective invention by Allen. The basic argument here is that free revealing of quality code can increase a programmer's reputation with peers (Raymond 1999) as well as his or her value on the job market (Lerner and Tirole 2002). Other arguments focus on the likely low level of damage to the innovator associated with a decision to freely reveal. First, with regard to the cost disadvantage incurred by innovators relative to free riders, consider that those who do contribute code to open source projects report that they both enjoyed and learned from the work of coding – thus the actual cost of generating that code may be even seen as negative by the writer (Lakhani and Wolf 2002). Second, with respect to the potential avoidability of losses associated with free revealing, consider that code-writers may often feel that others have developed code similar to their own and are likely to freely reveal it if they don't – which would render any decision that they might make to hide their code moot (Lakhani and von Hippel 2000). Finally, a number of writers have proposed that communal

norms, including altruism, may play a strong role in inducing free revealing in the field of open source field. For example, programmers may feel incited by “generalized reciprocity” (Eckh 1974) to reveal their code because they have benefited from the code freely revealed by others (e.g., Raymond 1999).

We propose an additional class of incentives for free revealing that are consequences of increased diffusion. When an innovating user freely reveals an innovation, the direct result is to increase the diffusion of that innovation relative to what it would be if the innovation were either licensed at a fee or held secret. The innovating user may then benefit from the increase in diffusion via a number of effects. Among these are network effects, reputational gains, and related innovations induced among and revealed by other users. In addition, and very importantly, an innovation that is freely revealed and adopted by others can become an informal standard that may preempt the development and/or commercialization of other versions of the innovation. If, as was suggested by Allen, the innovation that is revealed is designed in a way that is especially appropriate to conditions unique to the innovator, this can result in creating a permanent source of advantage for that innovator. Note that being first to reveal a given type of innovation increases a user’s chances of having *his* innovation widely adopted, other things being equal. This may induce innovators to race to reveal first.

Negative effects from free revealing are associated with any advantage this action provides to free-riding competitors relative to the innovator. The stronger the competition between the user-innovator and other users of the innovation, the larger will be the loss of competitive advantage that the innovator incurs by revealing. When competition between innovation users is low, e.g. due to geographical separation of markets, the revealing user does not suffer as a consequence of the advantages he provides to others. As was noted above, this loss is moderated by the degree to which the innovation suits other users’ needs: the more specific it is to the innovating user, the smaller the effect of competitors’ using it.

2.2 Qualitative evidence for free revealing

2.2.1 Equipment to produce “copper-interconnect” semiconductors

IBM was first to develop a process to manufacture semiconductors that incorporated copper interconnections among circuit elements instead of the traditionally-used aluminum ones. This innovation provided a major improvement to semiconductor performance, and on the face of it, it would have paid IBM to not reveal its process to others. After a delay IBM did, however, “freely reveal” increasing amounts of proprietary process information to rival users and to equipment suppliers.

IBM freely revealed information about its innovation because it needed equipment to implement the process on a production scale. Detailed design and production of such equipment required the combining of information held by semiconductor equipment suppliers and IBM. Since development of novel process equipment is a very expensive matter in the semiconductor field, suppliers would only be willing to build equipment that could potentially be sold to the entire market. Revealing innovation-related information selectively so as to insure that the IBM process approach became the industry standard – but that IBM still maintained a lead in the marketplace – was therefore in IBM’s best interest. Thus, IBM was motivated to “openly reveal” its innovation by the incentive of inducing manufacturer improvements and of setting a standard advantageous to the innovator user (Lim 2000, Harhoff 1996²).

2.2.2 Improvements to clinical chemistry analyzer equipment and tests

The Technicon Corporation was the first to produce automated clinical chemistry analyzers – a type of medical equipment used to determine the levels of chemical constituents of blood. The basic design of their product was taken from a system that had been earlier developed and used by laboratory clinicians – users of that type of equipment. This user-developed design was modular and well-suited to low-cost modification by other users who had an incentive to do this. After commercial introduction of the basic analyzer, many users developed test and hardware innovations and freely revealed these via publication and other means. Other users and the Technicon company then adopted many of these innovations (20+ in the case of the Technicon company itself) without payment to the innovators. Open revealing of innovations by users was encouraged by Technicon via a firm-supported research publication and via a research seminar series sponsored by the company (von Hippel and Finkelstein 1979).

Open revealing of innovations by users in this field served the purpose of inducing manufacturer improvements. Innovating users in this instance were typically employees in the clinical labs of publicly-supported institutions. They were scientist-rivals rather than commercial rivals with respect to other users. Their ongoing research benefited when Technicon adopted and improved their innovations. They also gained reputa-

2 Lim (2000) describes IBM’s motivation as follows: “IBM is also attempting to capture indirect benefits. According to my interviews, several IBM employees realized in the mid-1990s that it would benefit from lower equipment costs if the rest of the industry also adopted copper technology. This is consistent with the strategic sharing of information (Harhoff 1996). In line with this, IBM formed the alliance with Novellus and later relaxed somewhat on its secrecy. Although it continues to guard sensitive process information, it has begun to share its copper technology with other companies, including Siemens (Infineon), Sanyo, and a startup foundry in Taiwan (Table 8a). However, it is important to point out that this strategy of sharing technology was only feasible after IBM had established itself as the leader. Otherwise, another firm might have exploited the knowledge to beat it to market.” See Lim (2000) for a detailed discussion.

tion-related benefits from peers and employers when they established their priority with respect to their innovations.

2.2.3 Improvements to computerized library information systems

Library „OPACs“ (Online Public Access) are computerized information systems that give patrons access to library collections (they functionally replace the traditional „card catalog“ form of collection index) and to the rich information resources of the Internet. The first OPACs were developed by leading-edge users in the early 1970’s. In the late 1970’s, suppliers began to offer OPACs as commercial products. Prior to the late 1990’s OPACs were not designed to make modification by users easy. Nonetheless, a study of Australian libraries using OPACs showed that 26% had modified the code of the OPAC they had purchased (or developed themselves) to improve its functionality. The study also showed that innovating users freely revealed their innovations to other users and to their OPAC suppliers in manufacturer-sponsored “user’s group” meetings and elsewhere. Manufacturers were willing to adopt user-requested or user-prototyped improvements in OPAC functionality – without any payment to user-innovators - if “enough” users wanted the same thing. So one or more users would sometimes engage in pre-meeting lobbying to generate widespread support for an innovation they wished to have adopted and supported by their OPAC supplier. Innovating users in the library field were not rivals in the marketplace – each served a geographically or topically distinct group of patrons. They benefited from openly revealing their innovations by inducing manufacturer improvements (Morrison et al. 2000).

2.2.4 Development of open-source software

Open-source software products, such as Linux operating system software and Apache computer server software, are built up from modules and “patches” developed by programmer-users of that software. Open source software has its roots in the “free software” movement started by Richard Stallman in the early 1980s. Stallman founded the Free Software Foundation (FSF) as a means to counter the trend towards proprietary development of software packages, and the release of software without the underlying source code.³ Stallman’s pioneering idea was that software authors interested in pre-

3 Source code is a sequence of instructions to be executed by a computer to accomplish a program’s purpose. Programmers write computer software in the form of source code, and also “document” that source code with brief written explanations of the purpose and design of each section of their program. To convert a program into a form that can actually operate a computer, source code is translated into machine code using a software tool called a compiler. The compiling process removes program documentation and creates a “binary” version of the program - a sequence of computer instructions consisting only of strings of ones and zeros. Binary code is very difficult for programmers to read and interpret. Therefore, programmers or firms that wish to prevent others from understand-

servicing the status of their software as “free” software could use their own copyright to grant licenses on terms that would guarantee a number of rights to all future users. They could do this by simply affixing a standard license to their software that conveyed these rights. The basic license developed by Stallman to implement this idea was the General Public License or GPL (sometimes referred to as “copyleft” - a play on the word “copyright”). Basic rights transferred to those possessing a copy of free software include the right to use it at no cost, the right to study its “source code,” to modify it, and to distribute modified or unmodified versions to others at no cost.⁴

The philosophy of the FSF movement has been recently extended by a number of individuals promoting the “Open Source” concept. These individuals have been less concerned about the freeness of “free software” and have instead been interested in encouraging software companies to release source code for their products under the GPL or licenses conveying similar rights.

Contributors of open source code to an open source project freely reveal their developments to all other users and also to the volunteer individual or user organization that manages that open source software product. This individual or organization is responsible for approving software code and generally distributes improvements adopted as “official” by posting revised software code on the Internet. Innovating users receive no direct payment from freely revealing their innovations in this manner.

Students of the open source software development process report that innovating users have a number of motives for freely revealing their code to open source project managers and open source code users in general. If they freely reveal, others can debug and improve upon the modules they have contributed, to everyone’s benefit. They are also motivated to have their improvement incorporated into the standard version of the open-source software that is generally distributed by the volunteer open source user organization, because it will then be updated and maintained without further effort on the innovator’s part. This volunteer organization is the functional equivalent of a manufacturer with respect to inducing manufacturer improvements, because a user-developed improvement will *only* be assured of inclusion in new “official” software releases if it is approved and adopted by the coordinating user group. Innovating users also report be-

ing and modifying their code will release only binary versions of the software. In contrast, programmers or firms that wish to enable others to understand and update and modify their software will provide them with its source code.

4 Note that this is not full “free revealing” as defined in section 1, since others’ rights to use the software are restricted in some respects. Specifically, users of open source software protected under a GPL license may not incorporate that software into commercial, closed-source software. However, the rights granted to other users are so far-reaching that we judge that we may consider open source software as being freely revealed for the purposes of our analyses in this paper.

ing motivated to freely reveal their code under a free or open source license by a number of additional factors. These include giving support to open code, “giving back” to those whose freely-revealed code has been of value to them and gaining a personal reputational advantage by having a contribution accepted into the official version of an open source program (Lakhani and Wolf 2001).

3 A model of the relative benefits of secrecy and free revealing

It will be useful to explore the conditions under which secrecy or free revealing would be the most profitable course of action for users that have developed an innovation. We develop a model that allows us to systematically assess the likely profitability of users’ decisions to voluntarily reveal or hide information about their proprietary innovations. Our model incorporates four important variables related to this decision: (1) the intensity of competition among users (model parameter α); (2) the degree to which the innovation has a bias favoring the innovating user (parameter γ)⁵; (3) the value to the innovating user of the improvements that free revealing induces the manufacturer of the innovation to make and distribute to all users equally (parameter μ); and (4) the cost to each user of adopting the improved commercial product (parameter c).

In order to make our model more concise and also conservative with respect to the value of free revealing, we do not include the effect of two incentives mentioned in our literature review as likely to contribute to the profitability of freely revealing innovation-related information. These are: (1) the impact of community expectations favoring free revealing such as generalized reciprocity; and (2) any increase in the value of assets complementary to the innovation held by the firm – assets ranging from physical assets to reputational assets. For simplicity, we also do not consider any costs an innovator may incur to either maintain secrecy or increase the diffusion of freely-revealed information.

We also do not incorporate a number of related issues that may be of interest for further study. For example, we do not model explicitly how the heterogeneity among product users emerges, i.e., why some of them are endowed with innovations in our model and others are not. For our purposes here, it is immaterial whether innovative users have reached their position by chance or, for example, by a special ability to develop an innovative product or process. We also do not model explicitly why transactions are not accompanied by monetary compensation – for example, why the revelation

5 Innovations developed by users will typically, intentionally or unintentionally, contain imbedded characteristics particularly beneficial to that user. Manufacturers may wish to strip these out in order to make the innovation more general. However, it is frequently the case that manufacturers cannot identify all user-specific features embedded in a product or service, because the innovating user typically has a much deeper understanding of the innovation’s intended use and use context.

of an innovation is not subject to a licensing contract. We do this because, as was mentioned earlier, empirical evidence clearly indicates that this behavior is rare, and that the practical choice offered to user-innovators is the choice between enjoying some period of exclusivity via secrecy and the voluntary free revealing of their innovations without monetary compensation.

3.1 Model structure

We consider the case of two user firms each of which may have developed an innovation. We will assume that these user firms are identical with respect to all other aspects that bear on their profitability. The innovating firm(s) can profit from the innovation by keeping it secret (Mansfield 1985, Levin et al. 1987) or by revealing it to a manufacturer firm which will decide whether to improve the product and offer it to both users. Revealing the innovation to the manufacturer also entails making it known to the world at large. We make the conservative assumption that it is not possible to reveal the innovation exclusively to the manufacturer. One reason for this is that the manufacturer may be able to assess the profitability of incorporating the innovation into his product only after the market has had a chance to get to know it. The users may then decide to adopt or not to adopt the improved innovation. These decisions are modeled in a sequential-move game.

The basic parameters of the game are four. First, the extent of competition (denoted by α) between the user firms relates increases in the payoff of one firm to losses for the other one. The higher the payoff of firm 1, the lower will be the payoff of firm 2 if competition is intense. Conversely, if competition is low (due to largely separate markets, e.g.) then the impact of firms' payoffs on each other will be small.

Second, we allow for different technologies employed by the two user firms. While standard oligopoly models typically assume that firms are identical with respect to their production technology, we find this assumption quite restrictive and in many cases unrealistic. The usual logic would suggest that firms will adopt the lowest-cost technology and that in the long run, only firms with that technology will be present in the market. In reality, the choice of production technology is likely to be path-dependent. Due to sunk cost investments, firms will not adjust their production technology immediately, and even over longer periods of time we are likely to observe pronounced differences in production technology. The most important implication of heterogeneity of this kind is that the value of innovations may differ across firms. In particular, if firm 1 develops an innovation, it is likely to be tailored to its own production technology. Transferred to a different production environment, the beneficial effect of the innovation (e.g. its contribution to cost reduction) may be substantially lower than in the innovator's context. We capture the idiosyncratic nature of the innovation in a sepa-

rate parameter γ (“generality”) which ranges between 0 (the technology is completely specific and cannot usefully be employed by any other firm) and 1 (the technology is completely general and yields identical benefits to all firms).

The third parameter μ describes the extent to which the manufacturer can improve the innovation or reduce its price, e.g. by virtue of large-scale production. To the innovator, the value of the improved innovation is by a factor of $(1+\mu)$ larger than that of the “homebuilt” version. Naturally, in an extended model this would again be considered an endogenous variable, but for the time being we consider it as exogenously given.

Finally, when the manufacturer sells the improved innovation, adoption is associated with some costs c which reflect either switching costs or the price that the manufacturer demands for the improved product. We do not model the pricing decision; but we assume that the manufacturer will not be able to extract all the surplus from the transaction.

3.2 Case 1: Innovation by one user

Consider the case of user 1 having developed an innovation and user 2 not having done so – the game tree and payoffs for this situation are depicted in Figure 1.

Include Figure 1 about here

The undisclosed innovation yields an increment δ in present discounted profits to the innovator. But user 1’s gain also has a competitive effect for user 2. We note the strength of competition by α and specify the impact of the innovation on the other user’s profit as $-\alpha\delta$, where $0 \leq \alpha \leq 1$. In a fully developed oligopoly model, α would be a function of technical and economic determinants.⁶ Note that a more dramatic improve-

⁶ A commonly used approach models the demand that firm 1 faces as

$$D_1(p_1, p_2, a) = 1 - p_1 + ap_2,$$

where p_1 and p_2 denote the prices of firms 1 and 2, resp., and a parameterizes the degree of competition. Ignoring fixed cost for simplicity, the profit of firm 1 obtains as

$$\Pi_1(p_1, p_2, c_1, a) = (p_1 - c_1)(1 - p_1 + ap_2),$$

with c_i denoting the variable cost per unit of firm i . A calculation of Nash equilibrium strategies and profits shows that the variation of profits Π_1 and Π_2 with e.g. c_1 can be reasonably well approximated, in the relevant parameter range, by a linear function. Obviously, with a decrease in c_1 – due to an innovation by firm 1, say – Π_1 increases, while Π_2 decreases. Because of the near-linearity of the profit functions, this profit reduction for firm 2 is approximately proportional to the gain for firm 1. It is more pronounced the stronger competition is, i.e., the higher the competition parameter a . This

ment of the innovator's position is likely to hurt its competitors more than a marginal improvement would.

In stage 1, the innovative user 1 may decide to reveal the innovation to a manufacturer of the improved good. We assume that this revelation automatically triggers a complete spillover to the second user. Revealing may have the advantage that the manufacturer may have specific expertise in improving the product further which the innovator may lack. Moreover, the manufacturer may be able to produce the product at lower cost than the innovator himself would incur to reproduce it. But revealing has two disadvantages: i) the revealed information becomes available to all other agents, and ii) the manufacturer will offer improved or lower-cost versions of the innovation to all users on an equal basis. Thus, the payoffs in the second-most left branch in Figure 1 change accordingly.

At stage 2, the manufacturer decides whether to incorporate the innovation in his product and whether to spend effort on improvements. We assume that the manufacturer will do so only if the improved innovation can be sold to both users.⁷ Once the improved innovation is adopted, its effect on user 1's profit is given by $\Delta=(1+\mu)\delta$ with $\mu\geq 0$.

The second user firm may also profit from the innovation, but to a lesser extent than the innovator. This effect is due to the fact that the innovator will have tailored the innovation optimally to its own production environment. This specificity – we assume – cannot be unraveled by the manufacturer. Hence, the other user will only enjoy a marginal direct payoff of $\gamma\Delta$ with $0\leq\gamma\leq 1$. The case of $\gamma=1$ denotes one of complete generality of the innovation, while $\gamma=0$ denotes the polar case of complete specificity.

The impact of competition enters the payoffs again by subtraction of the other user's payoff times the competition parameter α . Hence, once the innovation has been revealed by the innovator, and once it has been adopted by his competitor, the innovator's payoff is given by $\Delta-\alpha\gamma\Delta-c$ while the other firm enjoys a payoff of $\gamma\Delta-\alpha\Delta-c$, where c is the respective user's cost of adoption.

User 2 adopts if

$$\gamma\Delta - \alpha\Delta - c > \gamma\delta - \alpha\Delta \quad \Leftrightarrow \quad \gamma\mu > c/\delta . \quad (1)$$

shows that the payoffs we used can qualitatively also be deduced from a microeconomic model. The big advantage of our approach lies in the much simpler mathematical expressions.

7 More formally, but with the same effect, we could assume that the fixed costs of manufacturer R&D will only be covered if the manufacturer can sell the innovation to both users.

Note that user 1 always adopts the improved innovation if user 2 does so (for user 1, the condition corresponding to (1) reads $\mu > c/\delta$, which is, since $\gamma \leq 1$, fulfilled whenever (1) is). Hence, from the manufacturer's point of view, it is only user 2's decision that is of interest.

In Figure 2, which shows the parameter space (α, γ) , this condition means that the parameter constellation (α, γ) must lie above the straight horizontal line. Given that user 2 would adopt a manufacturer-improved innovation, user 1 finds it profitable to reveal his innovation if

$$\Delta - \alpha\gamma\Delta - c > \delta \Leftrightarrow \gamma < (\mu - c/\delta) / (\alpha(\mu + 1)) \quad (2)$$

In the six panels in Figure 2, this condition is fulfilled in the area below the downward-sloping curve. Taking conditions (1) and (2) together yields the condition that the parameters (α, γ) must lie in the shaded area.

Include Figure 2 about here

In Figure 2, we depict these conditions for six different parameter combinations. The upper three scenarios assume that $\mu=1$, i.e., the manufacturer-improved innovation has twice the value to the innovator than the self-developed one. In the lower three panels, we assume that $\mu=2$. In the left-hand panels, the cost of adoption are relatively small relative to the innovation's value ($c/\delta = 1/4$). The middle and right-hand scenarios then display the free-revealing conditions for relative adoption cost levels of $c/\delta = 1/2$ and $c/\delta = 3/4$, respectively. The changes in μ and in c/δ affect both curves simultaneously. Higher adoption costs make adoption by user 2 less likely, and they also reduce the benefit from free revealing to the innovator. Similarly, an increase in the degree of improvement μ by the manufacturer induces greater willingness to adopt (user 2) and to reveal (user 1). Note that greater competition will tend to reduce the likelihood of free-revealing through its impact on the expected benefit for the innovator.

3.3 Case 2: Innovation by both users

If both users have innovated, then either none, one, or both of them may reveal their innovation to the manufacturer. Since such a revelation is observable by the competitor only with a considerable delay, we model the users' decisions to reveal as a simultaneous one.

Given that at least one innovator has revealed his innovation, the manufacturer has to decide if to incorporate it into his product or not. If he does, then in the following stage both users are facing the decision to adopt the improved product or not. Figure 3 shows the payoff matrices for the case that only one user has revealed (left-hand panel -

we depict the case that this is user 1) and for the case that both have revealed (right-hand panel).

Include Figure 3 about here

We assume that a user-innovator who does not adopt a manufacturer-improved product exclusively benefits from his own innovation (δ), not from his competitor's revealed innovation. The latter might yield him a benefit of $\gamma\delta$, but as we consider the two users' innovations as close substitutes, there is no point in adopting the other's innovation on top (and certainly not instead) of one's own. If a user who has not revealed his innovation adopts the manufacturer-improved product (which builds upon his competitor's innovation), we assume that the user's own innovation does not convey any additional benefits. This is a sensible assumption since, apart from their idiosyncrasies, the users' innovations are substitutes (and the revealing user's idiosyncrasy is built into the product).

If only user 1 has revealed, then adopting the manufacturer-improved product pays for him if

$$\Delta - c > \delta \Leftrightarrow \mu > c/\delta . \quad (3)$$

For the non-revealing user 2, adopting the innovation of user 1 is profitable if

$$\gamma\Delta - c > \delta \Leftrightarrow \gamma(1+\mu) - 1 > c/\delta . \quad (4)$$

Note that both conditions are independent of the respective competitor's action: the game has dominant strategies. Inequality (4) is stronger than (3). This means that three cases can be distinguished: if (4) is satisfied, then both users adopt the improved product. If inequality (3) holds, but not (4), then only user 1 (who had revealed his innovation) would adopt. This makes sense, since, due to specificity to user 1's needs, user 1 derives a higher utility from the improved product than user 2. If also (3) is violated, then the cost of adoption is too high even for user 1.

If both have revealed, then the manufacturer has to choose one version of the innovation to be implemented in his improved product. We do not model this choice in detail, but simply assume that the payoffs are now given by $\gamma^M\Delta - \alpha\gamma^M\Delta$ for both firms, where $\gamma \leq \gamma^M \leq 1$. This inequality would be consistent with a random choice by the manufacturer which leaves the expected payoff to both users somewhere in between the benefit of the single innovative user after revelation and the non-innovator after adoption. It can also be interpreted as a conscious choice by a manufacturer which – by inspection of two alternative innovations – gains some insight into how to combine design fea-

tures. This leads us to a symmetric final-stage payoff matrix (Figure 3, right-hand panel), where the condition for adoption of the innovation is the same for both players:

$$\gamma^M \Delta - c > \delta \quad \Leftrightarrow \quad \gamma^M(1+\mu) - 1 > c/\delta . \quad (5)$$

Since $\gamma \leq \gamma^M \leq 1$, the three conditions (3) to (5) lead us to distinguish four different cases, as illustrated in Figure 4.

Include Figure 4 about here

If we assume that incorporating the revealed innovation(s) into a product only pays for the manufacturer if both users adopt the improved product, then of the four cases (i) to (iv), only (i) and (ii) matter: in case (iii), only the revealing innovator (in case only one of them reveals) would adopt, and in case (iv) no-one would adopt.

Using the subgame payoff matrices depicted in Figure 3, one can construct reduced payoff matrices for the users' simultaneous decisions to reveal or not. These are shown in Figure 5.

Include Figure 5 about here

In case (i), that is, with low cost of adoption, revealing by both user-innovators is always an equilibrium, since

$$\gamma^M \Delta(1-\alpha) - c > \gamma \Delta - \alpha \Delta - c \quad \Leftrightarrow \quad \gamma^M - \gamma^M \alpha > \gamma - \alpha \quad (6)$$

and $\gamma \leq \gamma^M \leq 1$.

There may exist a second equilibrium in which none of the innovators reveals. The condition for this equilibrium to exist is

$$\delta - \alpha \delta > \Delta - \alpha \gamma \Delta - c \quad \Leftrightarrow \quad \gamma < (\alpha + \mu - c/\delta) / (\alpha(1+\mu)) . \quad (7)$$

In case (ii), “not reveal” by both users is obviously an (indifferent) equilibrium. Revelation by both users is an equilibrium if

$$\gamma^M \Delta(1-\alpha) - c > \delta - \alpha \delta \quad \Leftrightarrow \quad (\gamma^M(1+\mu) - 1)(1-\alpha) > c/\delta . \quad (8)$$

This condition is stronger than condition (5), which defines the upper limit for c/δ in case (ii). Depending on the values of α , γ and γ^M , there may or may not be values of c/δ in the range of case (ii) which fulfill inequality (8).

Figure 6 illustrates the case of two innovators. The chosen parameter combinations of μ and c/δ are the same as in Figure 2. In order to simplify the presentation, we have set $\gamma^M=(1+\gamma)/2$. However, other values of γ^M (satisfying $\gamma \leq \gamma^M \leq 1$) would yield qualitatively the same picture.

Include Figure 6 about here

The upper area in each graph, the full gray rectangle (areas A, B, C in graph a)⁸, corresponds to case (i). That is, inequality (4) is fulfilled, and revealing (followed by adoption) by both user-innovators is an equilibrium. “Not revealing” by both innovators is another equilibrium above the downward sloping curve (C in graph a), but not below it (A, B). The broken line separating A and B is defined by inequality (8). While (8) has been derived for case (ii) (areas D and E in graph a), it has a very interesting interpretation also in case (i): above the line (area A), the unique equilibrium “revealing” by both is preferred to the alternative symmetric situation “not revealing”. In area B, on the other hand, the user-innovators would do collectively better by not revealing, but a prisoner’s dilemma does lead them to reveal. The interpretation is insightful: due to higher competition compared to area A, the gains from using the improved product are largely competed away, and the cost c of adoption further reduces the players’ payoffs. However, if only one user-innovator were to reveal, then the other would still find it preferable to adopt. Doing so, he prevents his competitor from attaining too much of an advantage. But, given that he will adopt even in that case, he does better by also revealing, since then the manufacturer-improved product fits his idiosyncratic needs just as well as those of his competitor.

In area C, again the two players’ payoffs are higher if no-one reveals than if both reveal. However, in C, “no-one reveals” is also an equilibrium, such that the prisoner’s dilemma situation from B has turned into a coordination problem. The intuition behind this is that with increasing generality of the innovations (higher γ) and increasing competition (higher α), the relative gains from being the only one to reveal diminish, which turns “no-one reveals” into a stable situation.

The rectangle below, which fills the whole rest of the parameter space in graphs d) and e), corresponds to case (ii), where inequality (5) holds. Here, revealing by both users is an equilibrium above the upward sloping curve (area D in graph a). In this equilibrium, both innovators do better than in any other combination of strategies (which invariably yield the payoff $\delta-\alpha\delta$ for both, since in case (ii) no innovation is built into a product if only one user adopts). As we would expect, this equilibrium disappears for

⁸ The same areas A to F, although shifted and deformed, can be identified in all other graphs b) to f), except for area F in graphs d) and e)

high degrees of competition (high α) and low generality of the innovations (small γ), since both effects decrease the relative gains of revealing compared to secrecy.

In the whole rectangle (D, E), as well as in the space below (F), “no-one reveals” is also an equilibrium. More precisely, it is an indifferent equilibrium: if no-one reveals, there is no point in a unilateral deviation because the other user would not adopt the improved product – due to higher idiosyncrasy (lower γ) compared to case (i), it would be too specific to his competitor’s needs. This induces the manufacturer not to build it in the first place, such that a revelation by only one innovator would be without effect. In real life, this equilibrium should even be a stable situation, not just indifferent, since there will be some level of cost connected with revealing. Note that the equilibrium “no-one reveals” in areas D, E, and F is of a different nature than the corresponding one in area C. In C, being the only one to reveal does lead the manufacturer to incorporate the innovation, and the competitor to adopt it. However, due to high competition and low idiosyncrasy, this is disadvantageous to the revealing innovator.

Comparing graphs a) to f) illustrates the influence of the manufacturer’s improvement (μ) and the relative cost (c/δ) of adopting the improved product. The more pronounced the manufacturer’s value addition, the more attractive revealing becomes. When, as in graph c), this value addition is more modest ($\mu=1$) and the relative cost of adopting is high ($c/\delta=3/4$), then revealing only happens if the innovations’ generality is very high. On the other hand, when the manufacturer makes considerable improvements and the cost of adoption is relatively low (graph d, $\mu=2$ and $c/\delta=1/4$), then free revealing occurs for almost all parameter combinations (α, γ): only in extreme cases of strong competition and high idiosyncrasy (white area bottom right) innovators prefer to keep secrecy. Finally, consider the situations where revealing by both is an equilibrium, but one which is inferior to “no-one reveals” (areas B, C). For extensive manufacturer improvements μ and low adoption cost c/δ (graph d), these situations are restricted to cases of strong competition and medium to high generality. On the other extreme, if μ is low and c/δ is high (graph c), then “inefficient revealing” (from the point of view of the innovators) only occurs for high generality, for all but the lowest levels of competition.

4 Discussion

Innovation is often a process to which several actors with complementary capabilities contribute. Creating the conditions for these actors to engage jointly in innovation processes is often welfare-improving, since none of them has sufficient knowledge or information to produce the innovation on their own. Conversely, factors that prevent actors from participating in these processes would tend to reduce social welfare. The usual candidates for factors having such effects are transactions costs, informational asymmetries and the incompleteness of contracts. Taken alone or together, they may prevent economic agents from coming to the division of labor that allows each actor to

most effectively contribute to the development of innovations. Moreover, the classical view includes the assumption that any uncompensated information transfer that occurs in such a system must be involuntary – knowledge spills out in this case, to the detriment of the party losing it to competitors or others, since each agent seeks to keep its information proprietary in order to reap the maximum possible return on it. The conventional prescription is often to strengthen the property rights that innovators can obtain.

We think that this view is limited, and that it may actually misrepresent a large number of real-world cases. We agree to the notion that important forces introduce friction, and that the market for intellectual property is imperfect. Our assumption that licensing is not feasible, and that there are no side payments between users and between innovator users and manufacturers reflects this view. But this strong assumption does not imply that there will be no information transfer. First, we have pointed out that information transfers which are not accompanied by monetary compensation are frequent. Moreover, they occur intentionally – hence, they can be called voluntary information spillovers.

Our paper has sought to contribute to the reconciliation of this actual economic behavior with economic theory. We developed a simplified game-theoretic model describing the interplay of incentives related to the decision to freely reveal for various intensities of rivalry in the marketplace. The incentives described in the model developed in section 3 have a good fit with the real-world conditions faced by user-innovators in the copper-interconnections, clinical chemistry analyzers, OPACs, and open source software innovations described in section 2. In each of these cases, innovations freely revealed by users were adopted by manufacturers who then made them available to all users via commercial sale. (Recall that, in the case of user innovations developed for open-source software products, the functional equivalent of the manufacturer is the volunteer user group having, by established custom, the exclusive right to add innovations to the “official” version of that product (Raymond 1999).)

Our model identifies a range of conditions under which users could be expected to benefit from freely revealing their innovations to others via manufacturers. Thus we propose that, in a world of self-interested agents with complementary capabilities, free revealing can be profitable. We further reason that free-revealing by innovating users is likely to be a *common* phenomenon. Consider that, for any area of innovation application, there are likely to be a number of users extant with relevant and valuable information, and who are experiencing different conditions with respect to the desirability of revealing that information. Note that all that is required for information diffusion to occur is that at least *one* such user be in a position to benefit sufficiently from openly revealing his version of the innovation-related information.

The model we have discussed in this paper does not directly address the possibility of direct user-to-user transfer without any involvement by a manufacturer or functional equivalent. However, direct user-to-user transfer does often occur, as has been documented in the case of process innovations in iron and steel (Allen 1983, von Hippel 1987, Schrader 1991). Our model is applicable to user-to-user transfers in which the receiving user creates an externality enjoyed by the sender. In this case, the parameter μ in our model can be interpreted as the size of the externality while the interpretation of the other parameters does not change. (Allen proposed that free revealing can be profitable for an innovating user when a revealed innovation is specific to an asset held by that innovator, and when revealing raises the value of that asset enough to offset the loss of profits associated with the revelation.)

Our work points to an interesting and possibly major problem in the empirical assessment of knowledge flows. In an illuminating survey of the spillover literature, Griliches (1994) has already provided a cautious view of the advances made in this field. With few exceptions, authors in the area of spillover measurement interpret the measured impact of “knowledge pools” on R&D expenditures or output as the impact of unintended knowledge flows. In the usual interpretation, therefore, these measures give us an idea of the importance of knowledge externalities.⁹ In contrast, our data and model suggest that the interpretation of spillover coefficients may not only measure the externality originating from unintended spillover flows, but also the impact of information that has been revealed intentionally. If the phenomenon of free-revealing of innovations by users is indeed interesting and important, then it will be valuable to extend and improve the model of the phenomenon that we have presented here, and to develop other types of models as well.

Finally, our argument that many users will have an incentive to freely diffuse information about their innovations makes measures to encourage and utilize innovations developed by users an attractive proposition from the viewpoint of the overall economy. If innovating users generally were not willing to freely reveal their innovations, then general diffusion could not occur unless either many users independently developed similar innovations, or a manufacturer either developed or licensed a similar innovation and offered it for general sale in the marketplace. Under such conditions it would clearly be more economical to encourage licensing by users and/or innovation by manufacturers – thus avoiding the cost of multiple independent innovations by users. Promoting the development, free revealing, and widespread utilization of user innovations may

9 Cohen and Levin (1989, 1990) have argued that this view may have to be amended: they argue that firms without the capability to absorb knowledge will not profit from external knowledge. For some empirical evidence supporting this view, see Harhoff (2000).

often be in the best interest of profit-seeking user-innovators, and welfare-improving as well.

References

Allen, R. C. (1983), "Collective Invention," *Journal of Economic Behavior and Organization*, 4, 1, 1-24.

Arora, A., A. Fosfuri, and A. Gambardella (2000), "Licensing in the Chemical Industry," Unpublished working paper, first version April 1998, this version September 2000, Carnegie Mellon University, Pittsburgh PA.

Bessen, S. M. and G. Maskin (2000), "Sequential Innovation, Patents, and Imitation," working paper 00-01, Sloan School, Massachusetts Institute of Technology.

Cohen, W. M. and D. A. Levinthal (1990), "The Implications of Spillovers for R&D Investment and Welfare: A New Perspective," *Administrative Science Quarterly*, 35, 128-152.

Cohen, W. M. and D. A. Levinthal (1989), "Innovation and Learning: The Two Faces of R&D," *Economic Journal*, 99, 569-596.

Cohen, W.M., Nelson, R.R., and Walsh, J.P. (2000) "Protecting their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or not)," NBER Working Paper 7522.

de Fraja, G. (1993), "Strategic Spillovers in Patent Races," *International Journal of Industrial Organization*, 11, 139-146.

Ekeh, P. P. (1974), *Social Exchange Theory: The Two Traditions*, Harvard University Press, Cambridge, MA.

Enos, J. L. (1962), *Petroleum Progress and Profits: A History of Process Innovation*, MIT Press, Cambridge, MA.

Franke, N. and S. Shah (2002) "How communities support innovative activities: an exploration of assistance and sharing among end-users," *Research Policy*, forthcoming.

Freeman, C. (1968), "Chemical Process Plant: Innovation and the World Market," *National Institute Economic Review*, 45, 2957.

Griliches, Z. (1994), "The Search for R&D Spillovers," *Scandinavian Journal of Economics*, 94 (Supplement), 29-47.

Harhoff, D. (1996), "Strategic Spillovers and Incentives for Research and Development," *Management Science*, 42, 6, 907-925.

Harhoff, D. (2000), "R&D Spillovers, Technological Proximity, and Productivity

- Growth - Evidence from German Panel Data,” *Schmalenbach Business Review*, 52, 3, 238-260.
- Hirschleifer, J. (1971) “The private and social value of information and the reward to inventive activity,” *American Economic Review* LXI, 4, 561-574.
- Knight, K. E. (1963), “A Study of Technological Innovation: The Evolution of Digital Computers,” unpublished Ph.D. dissertation, Carnegie Institute of Technology, Pittsburgh, PA.
- Lakhani, K. and E. von Hippel (2000), “Emerging Economics of Virtual Communities: “Free” User-To-User Assistance in Open Source Software,” MIT Sloan School of Management Working paper # 4117.
- Lakhani, K. and R. Wolf (2001), “Does Free Software Mean Free Labor? Characteristics of Participants in Open Source Communities,” BCG Survey Report, Boston, MA: Boston Consulting Group Report (available at <http://www.osdn.com/bcg/>).
- Lerner and Tirole (2002), “Some Simple Economics of Open Source,” *Journal of Industrial Economics*, forthcoming.
- Levin, R. C., A. Klevorick, R. R. Nelson and S. G. Winter (1987), “Appropriating the Returns from Industrial Research and Development,” *Brookings Papers on Economic Activity*, 3, 783-820.
- Levi-Strauss, C. (1949), English translation 1969, *The Elementary Structures of Kinship*, Beacon Press, Boston.
- Lilien, G., P. Morrison and E. von Hippel (2002), “Performance of lead user concept generation methods relative to traditional ones: An empirical study,” *Management Science*, forthcoming
- Lim, K. (2000), “The Many Faces of Absorbitive Capacity: Spillovers of Copper Interconnect Technology for Semiconductor Chips,” MIT Sloan School of Management Working paper # 4110.
- Lüthje, C. (2000), “Characteristics of innovating users,” Technical University of Hamburg-Harburg Working Paper # 8.
- Mansfield, E. (1986) “Patents and innovation: An empirical study,” *Management Science*, 32, 173-181.
- Mansfield, E. (1985), “How Rapidly Does New Industrial Technology Leak Out?,” *Journal of Industrial Economics*, 34, 217-223.
- Mansfield, E., Schwartz, M. and Wagner, S. (1981), “Imitation costs and patents: An empirical study,” *Economic Journal*, 91, 907-918.
- Merton, R. K. (1973), *The Sociology of Science: Theoretical and Empirical Investigations*, edited by N. W. Storer, University of Chicago Press, Chicago.

- Mishina, K. (1989), *Essays on Technological Evolution*, unpublished Ph.D. Thesis, Harvard University.
- Morin, R. (1993), "The Meaning of Freeware," *Unix Review*, 11, 12, 101-102.
- Morrison, P., J. Roberts and E. von Hippel (2000), "Determinants of User Innovation and Innovation Sharing in a Local Market," *Management Science*, 46, 1513-1527.
- Nelson, R. R. and S. G. Winter (1982), *An Evolutionary Theory of Economic Change*, Harvard University Press, Cambridge, MA.
- Ogawa, S. (1998), "Does Sticky Information Affect the Locus of Innovation? Evidence from the Japanese Convenience-Store Industry," *Research Policy*, 26, 777-790.
- Raymond, E. S. (1999), *The Cathedral and the Bazaar: Musings on Linux and Open Source by an Accidental Revolutionary*, O'Reilly, Cambridge, MA.
- Scherer, F.M., et al., (1959), *Patents and the corporation*, 2nd edn. Boston, privately published
- Schrader, S. (1991), "Informal Technology Transfer Between Firms: Cooperation Through Information Trading," *Research Policy*, 20, 153-170.
- Shah, S. (2000), "Sources and Patterns of Innovation in a Consumer Products Field: Innovations in Sporting Equipment," MIT Sloan School of Management Working paper # 4105 (March).
- Shaw, B. (1985), "The Role of the Interaction between the User and the Manufacturer in Medical Equipment Innovation," *R&D Management*, 15, 4, 283-92.
- Taylor, C. T. and Z. A. Silberston (1973), *The Economic Impact of the Patent System: A Study of the British Patent System*, Cambridge University Press, Cambridge, MA.
- Tirole, J. (1989), *The Theory of Industrial Organization*, MIT Press, Cambridge, MA.
- Urban, G. L. and E. von Hippel (1988), "Lead User Analyses for the Development of New Industrial Products," *Management Science*, 34, 5, 569-82.
- von Hippel, Eric (1987), "Cooperation Between Rivals: Informal Know-How Trading," *Research Policy*, 16, 291-302.
- von Hippel, Eric (1988), *The Sources of Innovation*, Oxford University Press, New York.
- von Hippel, E. and S. N. Finkelstein (1979), "Analysis of Innovation in Automated Clinical Chemistry Analyzers," *Science & Public Policy*, 6, 1, 24-37.
- von Hippel, E. and G. von Krogh (2002). "Exploring the Open Source Software Phenomenon: Issues for Organization Science," *Organisation Science*, forthcoming.

Wilson, R.W. (1975) “The Sale of Technology Through Licensing”, unpublished Ph.D. dissertation, Yale University, New Haven, Conn.