

6 | Why Users Often Freely Reveal Their Innovations

Products, services, and processes developed by users become more valuable to society if they are somehow diffused to others that can also benefit from them. If user innovations are not diffused, multiple users with very similar needs will have to invest to (re)develop very similar innovations, which would be a poor use of resources from the social welfare point of view. Empirical research shows that new and modified products developed by users often do diffuse widely—and they do this by an unexpected means: user-innovators themselves often voluntarily publicly reveal what they have developed for all to examine, imitate, or modify without any payment to the innovator.

In this chapter, I first review evidence that free revealing is frequent. Next, I discuss the case for free revealing from an innovators' perspective, and argue that it often can be the best *practical* route for users to increase profit from their innovations. Finally, I discuss the implications of free revealing for innovation theory.

Evidence of Free Revealing

When my colleagues and I say that an innovator “freely reveals” proprietary information, we mean that all intellectual property rights to that information are voluntarily given up by that innovator and all parties are given equal access to it—the information becomes a public good (Harhoff, Henkel, and von Hippel 2003). 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 subscription to a journal or for a field trip to an innovation site 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 possessor of the information does not profit from any such expenditures made by its adopters, the information itself is still freely revealed, according to our definition. This definition of free revealing is rather extreme in that revealing with some small constraints, as is sometimes done, would achieve largely the same economic effect. Still, it is useful to discover that innovations are often freely revealed even in terms of this stringent definition.

Routine and intentional free revealing among profit-seeking firms was first described by Allen (1983). He noticed the phenomenon, which he called collective invention, in historical records from the nineteenth-century English iron industry. In that industry, ore was processed into iron by means of large furnaces heated to very high temperatures. Two attributes of the furnaces used had been steadily improved during the period 1850–1875: chimney height had been increased and the temperature of the combustion air pumped into the furnace during operation had been raised. These two technical changes significantly and progressively improved the energy efficiency of iron production—a very important matter for producers. Allen noted the surprising fact that employees of competing firms publicly revealed information on their furnace design improvements and related performance data in meetings of professional societies and in published material.

After Allen's initial observation, a number of other authors searched for free revealing among profit-seeking firms and frequently found it. Nuvolari (2004) studied a topic and time similar to that studied by Allen and found a similar pattern of free revealing in the case of improvements made to steam engines used to pump out mines in the 1800s. At that time, mining activities were severely hampered by water that tended to flood into mines of any depth, and so an early and important application of steam engines was for the removal of water from mines. Nuvolari explored the technical history of steam engines used to drain copper and tin mines in England's Cornwall District. Here, patented steam engines developed by James Watt were widely deployed in the 1700s. After the expiration of the Watt patent, an engineer named Richard Trevithick developed a new type of high-pressure engine in 1812. Instead of patenting his invention, he made his

design available to all for use without charge. The engine soon became the basic design used in Cornwall. Many mine engineers improved Trevithick's design further and published what they had done in a monthly journal, *Leans Engine Reporter*. This journal had been founded by a group of mine managers with the explicit intention of aiding the rapid diffusion of best practices among these competing firms.

Free revealing has also been documented in the case of more recent industrial equipment innovations developed by users. Lim (2000) reports that IBM was first to develop a process to manufacture semiconductors that incorporated copper interconnections among circuit elements instead of the traditionally used aluminum ones. After some delay, IBM revealed increasing amounts of proprietary process information to rival users and to equipment suppliers. Widespread free revealing was also found in the case of automated clinical chemistry analyzers developed by the Technicon Corporation for use in medical diagnosis. After commercial introduction of the basic analyzer, many users developed major improvements to both the analyzer and to the clinical tests processed on that equipment. These users, generally medical personnel, freely revealed their improvements via publication, and at company-sponsored seminars (von Hippel and Finkelstein 1979). Mishina (1989) found free, or at least selective no-cost revealing in the lithographic equipment industry. He reported that innovating equipment users would sometimes reveal what they had done to machine manufacturers. Morrison, Roberts, and I, in our study of library IT search software (discussed in chapter 2 above), found that innovating users freely revealed 56 percent of the software modifications they had developed. Reasons given for not revealing the remainder had nothing to do with considerations of intellectual property protection. Rather, users who did not share said they had no convenient users' group forum for doing so, and/or they thought their innovation was too specialized to be of interest to others.

Innovating users of sports equipment also have been found to freely reveal their new products and product modifications. Franke and Shah (2003), in their study of four communities of serious sports enthusiasts described in chapter 2, found that innovating users uniformly agreed with the statement that they shared their innovation with their entire community free of charge—and strongly disagreed with the statement that they sold their innovations ($p < 0.001$, t-test for dependent samples). Interestingly, two of the four communities they studied engaged in activities involving significant

competition among community members. Innovators in these two communities reported high but significantly less willingness to share, as one might expect in view of the potentially higher level of competitive loss free revealing would entail.

Contributors to the many open source software projects extant (more than 83,000 were listed on SourceForge.net in 2004) also routinely make the new code they have written public. Well-known open source software products include the Linux operating system software and the Apache web server computer software. Some conditions are attached to open source code licensing to ensure that the code remains available to all as an information commons. Because of these added protections, open source code does not quite fit the definition of free revealing given earlier in this chapter. (The licensing of open source software will be discussed in detail in chapter 7.)

Henkel (2003) showed that free revealing is sometimes practiced by directly competing manufacturers. He studied manufacturers that were competitors and that had all built improvements and extensions to a type of software known as embedded Linux. (Such software is “embedded in” and used to operate equipment ranging from cameras to chemical plants.) He found that these manufacturers freely revealed improvements to the common software platform that they all shared and, with a lag, also revealed much of the equipment-specific code they had written.

The Practical Case for Free Revealing

The “private investment model” of innovation assumes that innovation will be supported by private investment if and as innovators can make attractive profits from doing so. In this model, any free revealing or uncompensated “spillover” of proprietary knowledge developed by private investment will reduce the innovator’s profits. It is therefore assumed that innovators will strive to avoid spillovers of innovation-related information. From the perspective of this model, then, free revealing is a major surprise: it seems to make no sense that innovators would intentionally give away information for free that they had invested money to develop.

In this subsection I offer an explanation for the puzzle by pointing out that free revealing is often the *best practical* option available to user innovators. Harhoff, Henkel, and von Hippel (2003) found that it is in practice very difficult for most innovators to protect their innovations from direct

or approximate imitation. This means that the practical choice is typically *not* the one posited by the private investment model: should innovators voluntarily freely reveal their innovations, or should they protect them? Instead, the real choice facing user innovators often is whether to voluntarily freely reveal or to arrive at the same end state, perhaps with a bit of a lag, via involuntary spillovers. The practical case for voluntary free revealing is further strengthened because it can be accomplished at low cost, and often yields private benefits to the innovators. When benefits from free revealing exceed the benefits that are *practically* obtainable from holding an innovation secret or licensing it, free revealing should be the preferred course of action for a profit-seeking firm or individual.

Others Often Know Something Close to “Your” Secret

Innovators seeking to protect innovations they have developed as their intellectual property must establish some kind of monopoly control over the innovation-related information. In practice, this can be done either by effectively hiding the information as a trade secret, or by getting effective legal protection by patents or copyrights. (Trademarks also fall under the heading of intellectual property, but we do not consider those here.) In addition, however, it must be the case that *others* do not know substitute information that skirts these protections and that they *are* willing to reveal. If multiple individuals or firms have substitutable information, they are likely to vary with respect to the competitive circumstances they face. A specific innovator’s ability to protect “its” innovation as proprietary property will then be determined for all holders of such information by the decision of the one having the least to lose by free revealing. If one or more information holders expect no loss or even a gain from a decision to freely reveal, then the secret will probably be revealed despite other innovators’ best efforts to avoid this fate.

Commonly, firms and individuals have information that would be valuable to those seeking to imitate a particular innovation. This is because innovators and imitators seldom need access to a specific version of an innovation. Indeed, engineers seldom even want to see a solution exactly as their competitors have designed it: specific circumstances differ even among close competitors, and solutions must in any case be adapted to each adopter’s precise circumstances. What an engineer does want to extract from the work of others is the principles and the general outline of

a possible improvement, rather than the easily redevelopable details. This information is likely to be available from many sources.

For example, suppose you are a system developer at a bank and you are tasked with improving in-house software for checking customers' credit online. On the face of it, it might seem that you would gain most by studying the details of the systems that competing banks have developed to handle that same task. It is certainly true that competing banks may face market conditions very similar to your bank, and they may well not want to reveal the valuable innovations they have developed to a competitor. However, the situation is still by no means bleak for an imitator. There are also many non-bank users of online credit checking systems in the world—probably millions. Some will have innovated and be willing to reveal what they have done, and some of these will have the information you need. The likelihood that the information you seek will be freely revealed by some individual or firm is further enhanced by the fact that your search for novel basic improvements may profitably extend far beyond the specific application of online credit checking. Other fields will also have information on *components* of the solution you need. For example, many applications in addition to online credit checking use software components designed to determine whether persons seeking information are authorized to receive it. Any can potentially be a provider of information for this element of your improved system.

A finding by Lakhani and von Hippel (2003) illustrates the possibility that many firms and individuals may have similar information. Lakhani and von Hippel studied Apache help-line websites. These sites enable users having problems with Apache software to post questions, and others to respond with answers. The authors asked those who provided answers how many other help-line participants they thought also knew a solution to specific and often obscure problems they had answered on the Apache online forum. Information providers generally were of the opinion that some or many other help-line participants also knew a solution, and could have provided an answer if they themselves had not done so (table 6.1).

Even in the unlikely event that a secret is held by one individual, that information holder will not find it easy to keep a secret for long. Mansfield (1985) studied 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

Table 6.1

Even very specialized information is often widely known. Tabulated here are answers to a question asked of help-line information providers: “How many others do you think knew the answer to the question you answered?”

	Frequent providers (<i>n</i> = 21)	Other providers (<i>n</i> = 67)
Many	38%	61%
A few with good Apache knowledge	38%	18%
A few with specific problem experience	24%	21%

Source: Lakhani and von Hippel 2003, table 10.

about a year.” This observation is supported by Allen’s previously mentioned study of free revealing in the nineteenth-century English iron industry. Allen (1983, p. 17) notes that developers of improved blast furnace designs were unlikely to be able to keep their valuable innovations secret because “in the case of blast furnaces and steelworks, the construction would have been done by contractors who would know the design.” Also, “the designs themselves were often created by consulting engineers who shifted from firm to firm.”

Low Ability to Profit from Patenting

Next, suppose that a single user-innovator is the only holder of a particular unit of innovation-related information, and that for some reason there are no easy substitutes. That user actually does have a real choice with respect to disposing of its intellectual property: it can keep the innovation secret and profit from in-house use only, it can license it, or it can choose to freely reveal the innovation. We have just seen that the practical likelihood of keeping a secret is low, especially when there are multiple potential providers of very similar secrets. But if one legally protects an innovation by means of a patent or a copyright, one need not keep an innovation secret in order to control it. Thus, a firm or an individual that freely reveals is forgoing any chance to get a profit via licensing of intellectual property for a fee. What, in practical terms, is the likelihood of succeeding at this and so of forgoing profit by choosing to freely reveal?

In most subject matters, the relevant form of legal protection for intellectual property is the patent, generally the “utility” patent. (The notable exception is the software industry, where material to be licensed is often

protected by copyright.) In the United States, utility patents may be granted for inventions related to composition of matter and/or a method and/or a use. They may not be granted for ideas per se, mathematical formulas, laws of nature, and anything repugnant to morals and public policy. Within subject matters potentially protectable by patent, protection will be granted only when the intellectual property claimed meets additional criteria of usefulness, novelty, and non-obviousness to those skilled in the relevant art. (The tests for whether these criteria have been met are based on judgement. When a low threshold is used, patents are easier to get, and vice-versa (Hall and Harhoff 2004).)

The real-world value of patent protection has been studied for more than 40 years. Various researchers have found that, with a few exceptions, innovators do *not* think that patents are very useful either for excluding imitators or for capturing royalties in most industries. (Fields generally cited as exceptions are pharmaceuticals, chemicals, and chemical processes, where patents do enable markets for technical information (Arora et al. 2001).) Most respondents also say that the availability of patent protection does not induce them to invest more in research and development than they would if patent protection did not exist. Taylor and Silberston (1973) reported that 24 of 32 firms said that only 5 percent or less of their R&D expenditures were dependent on the availability of patent protection. Levin et al. (1987) surveyed 650 R&D executives in 130 different industries and found that all except respondents from the chemical and pharmaceutical industries judged patents to be “relatively ineffective.” Similar findings have been reported by Mansfield (1968, 1985), by Cohen et al. (2000, 2002), by Arundel (2001), and by Sattler (2003).

Despite recent governmental efforts to strengthen patent enforcement, a comparison of survey results indicates only a modest increase between 1983 and 1994 in large firms’ evaluations of patents’ effectiveness in protecting innovations or promoting innovation investments. Of course, there are notable exceptions: some firms, including IBM and TI, report significant income from the licensing of their patented technologies.

Obtaining a patent typically costs thousands of dollars, and it can take years (Harhoff, Henkel, and von Hippel 2003). This makes patents especially impractical for many individual user-innovators, and also for small and medium-size firms of limited means. As a stark example, it is hard to imagine that an individual user who has developed an innovation in sports

equipment would find it appealing to invest in a patent and in follow-on efforts to find a licensee and enforce payment. The few that do attempt this, as Shah (2000) has shown, seldom gain any return from licensees as payment for their time and expenditures.

Copyright is a low-cost and immediate form of legal protection that applies to original writings and images ranging from software code to movies. Authors do not have to apply for copyright protection; it “follows the author’s pen across the page.” Licensing of copyrighted works is common, and it is widely practiced by commercial software firms. When one buys a copy of a non-custom software product, one is typically buying only a license to use the software, not buying the intellectual property itself. However, copyright protection is also limited in an important way. Only the specific original writing itself is protected, not the underlying invention or ideas. As a consequence, copyright protections can be circumvented. For example, those who wish to imitate the function of a copyrighted software program can do so by writing new software code to implement that function.

Given the above, we may conclude that in practice little profit is being sacrificed by many user-innovator firms or individuals that choose to forgo the possibility of legally protecting their innovations in favor of free revealing.

Positive Incentives for Free Revealing

As was noted earlier, 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 that all interested parties are given access to it—the information becomes a public good. These conditions can often be met at a very low cost. For example, an innovator can simply post information about the innovation on a website without publicity, so those potentially interested must discover it. Or a firm that has developed a novel process machine can agree to give a factory tour to any firm or individual that thinks to ask for one, without attempting to publicize the invention or the availability of such tours in any way. However, it is clear that many innovators go beyond basic, low-cost forms of free revealing. They spend significant money and time to ensure that their innovations are seen in a favorable light, and that information about them is effectively and widely diffused. Writers of computer code may work hard to eliminate all bugs and to document their code

in a way that is very easy for potential adopters to understand before freely revealing it. Plant owners may repaint their plant, announce the availability of tours at a general industry meeting, and then provide a free lunch for their visitors.

Innovators' *active* efforts to diffuse information about their innovations suggest that there are positive, private rewards to be obtained from free revealing. A number of authors have considered what these might be. Allen (1983) proposed that reputation gained for a firm or for its managers might offset a reduction in profits for the firm caused by free revealing. Raymond (1999) and Lerner and Tirole (2002) elaborated on this idea when explaining free revealing by contributors to open source software development projects. Free revealing of high-quality code, they noted, can increase a programmer's reputation with his peers. This benefit can lead to other benefits, such as an increase in the programmer's value on the job market. Allen has argued that free revealing might have effects that actually increase a firm's profits if the revealed innovation is to some degree specific to assets owned by the innovator (see also Hirschleifer 1971).

Free revealing may also increase an innovator's profit in other ways. 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. (The classic illustration of a network effect is that the value of each telephone goes up as more are sold, because the value of a phone is strongly affected by the number of others who can be contacted in the network.) 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 Allen suggested, 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.

Being first to reveal a certain type of innovation increases a user firm's chances of having its innovation widely adopted, other things being equal. This may induce innovators to race to reveal first. Firms engaged in a patent race may disclose information voluntarily if the profits from success do not go only to the winner of the race. If being second quickly is preferable to

being first relatively late, there will be an incentive for voluntary revealing in order to accelerate the race (de Fraja 1993).

Incentives to freely reveal have been most deeply explored in the specific case of open source software projects. 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 be assured of inclusion in new "official" software releases only if it is approved and adopted by the coordinating user group. Innovating users also report being 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 and "giving back" to those whose freely revealed code has been of value to them (Lakhani and Wolf 2005).

By freely revealing information about an innovative product or process, a user makes it possible for manufacturers to learn about that innovation. Manufacturers may then improve upon it and/or offer it at a price lower than users' in-house production costs (Harhoff et al. 2003). When the improved version is offered for sale to the general market, the original user-innovator (and other users) can buy 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 variation of this argument applies to the free revealing among competing manufacturers documented by Henkel (2003). Competing developers of embedded Linux systems were creating software that was specifically designed to run the hardware products of their specific clients. Each manufacturer could freely reveal this equipment-specific code without fear of

direct competitive repercussions: it was applicable mainly to specific products made by a manufacturer's client, and it was less valuable to others. At the same time, all would jointly benefit from free revealing of improvements to the underlying embedded Linux code base, upon which they all build their proprietary products. After all, the competitive advantages of all their products depended on this code base's being equal to or better than the proprietary software code used by other manufacturers of similar products. Additionally, Linux software was a complement to hardware that many of the manufacturers in Henkel's sample also sold. Improved Linux software would likely increase sales of their complementary hardware products. (Complement suppliers' incentives to innovate have been modeled by Harhoff (1996).)

Free Revealing and Reuse

Of course, free revealing is of value only if others (re)use what has been revealed. It can be difficult to track what visitors to an information commons take away and reuse, and there is as yet very little empirical information on this important matter. Valuable forms of reuse range from the gaining of general ideas of development paths to pursue or avoid to the adoption of specific designs. For example, those who download software code from an open source project repository can use it to learn about approaches to solving a particular software problem and/or they may reuse portions of the downloaded code by inserting it directly into a software program of their own. Von Krogh et al. (2004) studied the latter type of code reuse in open source software and found it very extensive. Indeed, they report that *most* of the lines of software code in the projects they studied were taken from the commons of other open source software projects and software libraries and reused.

In the case of academic publications, we see evidence that free revealing does increase reuse—a matter of great importance to academics. A citation is an indicator that information contained in an article has been reused: the article has been read by the citing author and found useful enough to draw to readers' attention. Recent empirical studies are finding that articles to which readers have open access—articles available for free download from an author's website, for example—are cited significantly more often than are equivalent articles that are available only from libraries or from

publishers' fee-based websites. Antelman (2004) finds an increase in citations ranging from 45 percent in philosophy to 91 percent in mathematics. She notes that "scholars in diverse disciplines are adopting open-access practices at a surprisingly high rate and are being rewarded for it, as reflected in [citations]."

Implications for Theory

We have seen that in practice free revealing may often be the best practical course of action for innovators. How can we tie these observations back to theory, and perhaps improve theory as a result? At present there are two major models that characterize how innovation gets rewarded. The private investment model is based on the assumption that innovation will be supported by private investors expecting to make a profit. To encourage private investment in innovation, society grants innovators some limited rights to the innovations they generate via patents, copyrights, and trade secrecy laws. These rights are intended to assist innovators in getting private returns from their innovation-related investments. At the same time, the monopoly control that society grants to innovators and the private profits they reap create a loss to society relative to the free and unfettered use by all of the knowledge that the innovators have created. Society elects to suffer this social loss in order to increase innovators' incentives to invest in the creation of new knowledge (Arrow 1962; Dam 1995).

The second major model for inducing innovation is termed the collective action model. It applies to the provision of public goods, where a public good is defined by its non-excludability and non-rivalry: if any user consumes it, it cannot be feasibly withheld from other users, and all consume it on the same terms (Olson 1967). The collective action model assumes that innovators are *required* to relinquish control of knowledge or other assets they have developed to a project and so make them a public good. This requirement enables collective action projects to avoid the social loss associated with the restricted access to knowledge of the private investment model. At the same time, it creates problems with respect to recruiting and motivating potential contributors. Since contributions to a collective action project are a public good, users of that good have the option of waiting for others to contribute and then free riding on what they have done (Olson 1967).

The literature on collective action deals with the problem of recruiting contributors to a task in a number of ways. Oliver and Marwell (1988) and Taylor and Singleton (1993) predict that the description of a project's goals and the nature of recruiting efforts should matter a great deal. Other researchers argue that the creation and deployment of selective incentives for contributors is essential to the success of collective action projects. For example, projects may grant special credentials to especially productive project members (Friedman and McAdam 1992; Oliver 1980). The importance of selective incentives suggests that small groups will be most successful at executing collective action projects. In small groups, selective incentives can be carefully tailored for each group member and individual contributions can be more effectively monitored (Olson 1967; Ostrom 1998).

Interestingly, successful open source software projects do not appear to follow any of the guidelines for successful collective action projects just described. With respect to project recruitment, goal statements provided by successful open source software projects vary from technical and narrow to ideological and broad, and from precise to vague and emergent (for examples, see goal statements posted by projects hosted on Sourceforge.net).¹ Further, such projects may engage in no active recruiting beyond simply posting their intended goals and access address on a general public website customarily used for this purpose (for examples, see the Freshmeat.net website). Also, projects have shown by example that they can be successful even if large groups—perhaps thousands—of contributors are involved. Finally, open source software projects seem to expend no effort to discourage free riding. Anyone is free to download code or seek help from project websites, and no apparent form of moral pressure is applied to make a compensating contribution (e.g., “If you benefit from this code, please also contribute . . .”).

What can explain these deviations from expected practice? What, in other words, can explain free revealing of privately funded innovations and enthusiastic participation in projects to produce a public good? From the theoretical perspective, Georg von Krogh and I think the answer involves revisiting and easing some of the basic assumptions and constraints conventionally applied to the private investment and collective action models of innovation. Both, in an effort to offer “clean” and simple models for research, have excluded from consideration a very rich and fertile middle

ground where incentives for private investment and collective action can coexist, and where a “private-collective” innovation model can flourish. More specifically, a private-collective model of innovation occupies the middle ground between the private investment model and the collective action model by:

- Eliminating the assumption in private investment models that free revealing of innovations developed with private funds will represent a loss of private profit for the innovator and so will not be engaged in voluntarily. Instead the private-collective model proposes that under common conditions free revealing of proprietary innovations may increase rather than decrease innovators’ private profit.
- Eliminating the assumption in collective action models that a free rider obtains benefits from the completed public good that are equal to those a contributor obtains. Instead, the private-collective model proposes that contributors to a public good can *inherently* obtain greater private benefits than free riders. These provide incentives for participation in collective action projects that need not be managed by project personnel (von Hippel and von Krogh 2003).

In summation: Innovations developed at private cost are often revealed freely, and this behavior makes economic sense for participants under commonly encountered conditions. A private-collective model of innovation incentives can explain why and when knowledge created by private funding may be offered freely to all. When the conditions are met, society appears to have the best of both worlds—new knowledge is created by private funding and then freely revealed to all.

