

8 | Adapting Policy to User Innovation

Government policy makers generally wish to encourage activities that increase social welfare, and to discourage activities that reduce it. Therefore, it is important to ask about the social welfare effects of innovation by users. Henkel and von Hippel (2005) explored this matter and concluded that social welfare is likely to be higher in a world in which both users and manufacturers innovate than in a world in which only manufacturers innovate.

In this chapter, I first explain that innovation by users complements manufacturer innovation and can also be a source of success-enhancing new product ideas for manufacturers. Next, I note that innovation by users does not exhibit several welfare-reducing effects associated with innovation by manufacturers. Finally, I evaluate the effects of public policies on user innovation, and suggest modifications to those that—typically unintentionally—discriminate against innovation by users.

Social Welfare Effects of User Innovation

Social welfare functions are used in welfare economics to provide a measure of the material welfare of society, using economic variables as inputs. A social welfare function can be designed to express many social goals, ranging from population life expectancies to income distributions. Much of the literature on product diversity, innovation, and social welfare evaluates the impact of economic phenomena and policy on social welfare from the perspective of total income of a society without regard to how that income is distributed. We will take that viewpoint here.

User Innovation Improves Manufacturers' Success Rates

It is striking that most new products developed and introduced to the market by manufacturers are commercial failures. Mansfield and Wagner

(1975) found the overall probability of success for new industrial products to be only 27 percent. Elrod and Kelman (1987) found an overall probability of success of 26 percent for consumer products. Balachandra and Friar (1997), Poolton and Barclay (1998), and Redmond (1995) found similarly high failure rates in new products commercialized. Although there clearly is some recycling of knowledge from failed projects to successful ones, much of the investment in product development is highly specific. This high failure rate therefore represents a huge inefficiency in the conversion of R&D investment to useful output, and a corresponding reduction in social welfare.

Research indicates that the major reason for the commercial failure of manufacturer-developed products is poor understanding of users' needs by manufacturer-innovators. The landmark SAPPHO study showed this in a very clear and convincing way. This study was based on a sample of 31 product pairs. Members of each pair were selected to address the same function and market. (For example, one pair consisted of two "roundness meters," each developed by a separate company.) One member of each pair was a commercial success (which showed that there *was* a market for the product type); the other was a commercial failure. The development process for each successful and failing product was then studied in detail. The primary factor found to distinguish success from failure was that a deeper understanding of the market and the need was associated with successful projects (Achilladelis et al. 1971; Rothwell et al. 1974). A study by Mansfield and Wagner (1975) came to the same conclusion. More recent studies of information stickiness and the resulting asymmetries of information held by users and manufacturers, discussed in chapter 3, support the reasonableness of this general finding. Users are the generators of information regarding their needs. The decline in accuracy and completeness of need information after transfer from user to manufacturer is likely to be substantial because important elements of this information are likely to be sticky (von Hippel 1994; Ogawa 1998).

Innovations developed by users can improve manufacturers' information on users' needs and so improve their new product introduction success rates. Recall from previous chapters that innovation by users is concentrated among lead users. These lead users tend, as we have seen, to develop functionally novel products and product modifications addressing their own needs at the leading edge of markets where potential sales are both

small and uncertain. Manufacturers, in contrast, have poorer information on users' needs and use contexts, and will prefer to manufacture innovations for larger, more certain markets. In the short term, therefore, user innovations will tend to *complement* rather than substitute for products developed by manufacturers. In the longer term, the market as a whole catches up to the needs that motivated the lead user developments, and manufacturers will begin to find production of similar innovations to be commercially attractive. At that point, innovations by lead users can provide very useful information to manufacturers that they would not otherwise have.

As lead users develop and test their solutions in their own use environments, they learn more about the real nature of their needs. They then often freely reveal information about their innovations. Other users then may adopt the innovations, comment on them, modify and improve them, and freely reveal what they have done in turn. All of this freely revealed activity by lead users offers manufacturers a great deal of useful information about both needs embodied in solutions and about markets. Given access to a user-developed prototype, manufacturers no longer need to understand users' needs very accurately and richly. Instead they have the much easier task of replicating the function of user prototypes that users have already demonstrated are responsive to their needs. For example, a manufacturer seeking to commercialize a new type of surgical equipment and coming upon prototype equipment developed by surgeons need not understand precisely why the innovators want this product or even precisely how it is used; the manufacturer need only understand that many surgeons appear willing to pay for it and then reproduce the important features of the user-developed prototypes in a commercial product.

Observation of innovation by lead users and adoption by follow-on users also can give manufacturers a better understanding of the size of the potential market. Projections of product sales have been shown to be much more accurate when they are based on actual customer behavior than when they are based on potential buyers' pre-use expectations. Monitoring of field use of user-built prototypes and of their adoption by other users can give manufacturers rich data on precisely these matters and so should improve manufacturer's commercial success. In net, user innovation helps to reduce information asymmetries between users and manufacturers and so increases the efficiency of the innovation process.

User Innovation and Provisioning Biases

The economic literature on the impact of innovation on social welfare generally seeks to understand effects that might induce society to create too many product variations (overprovisioning) or too few (underprovisioning) from the viewpoint of net social economic income (Chamberlin 1950). Greater variety of products available for purchase is assumed to be desirable, in that it enables consumers to get more precisely what they want and/or to own a more diverse array of products. However, increased product diversity comes at a cost: smaller quantities of each product will be produced on average. This in turn means that development-related and production-related economies of scale are likely to be less. The basic tradeoff between variety and cost is what creates the possibility of overprovisioning or underprovisioning product variety. Innovations such as flexible manufacturing may reduce fixed costs associated with increased diversity and so shift the optimal degree of diversity upward. Nonetheless, the conflict still persists.

Henkel and I studied the welfare impact of adding users as a source of innovation to existing analyses of product diversity, innovation, and social welfare. Existing models uniformly contained the assumption that new products and services were supplied to the economy by manufacturers only. We found that the addition of innovation by users to these analyses largely avoids the welfare-reducing biases that had been identified. For example, consider “business stealing” (Spence 1976). This term refers to the fact that commercial manufacturers benefit by diverting business from their competitors. Since they do not take this negative externality into account, their private gain from introducing new products exceeds society’s total gain, tilting the balance toward overprovision of variety. In contrast, a freely revealed user innovation may also reduce incumbents’ business, but not to the innovator’s benefit. Hence, innovation incentives are not socially excessive.

Freely revealed innovations by users are also likely to reduce deadweight loss caused by pricing of products above their marginal costs. (Deadweight loss is a reduction in social welfare that occurs when goods are sold at a price above their marginal cost of production.) When users make information about their innovations available for free, and if the marginal cost of revealing that information is zero, an imitator only has to bear the cost of adoption. This is statically efficient. The availability of free user innovations can also induce sellers of competing commercial offerings to reduce their prices, thus indirectly leading to another reduction in dead-weight loss.

Reducing prices toward marginal costs can also reduce incentives to over-provision variety (Tirole 1988).

Henkel and I also explored a few special situations where social welfare might be *reduced* by the availability of freely revealed user innovations. One of these was the effect of reduced pricing power on manufacturers that create “platform” products. Often, a manufacturer of such a product will want to sell the platform—a razor, an ink-jet printer, a video-game player—at a low margin or a loss, and then price necessary add-ons (razor blades, ink cartridges, video games) at a much higher margin. If the possibility of freely revealed add-ons developed by users makes development of a platform unprofitable for a manufacturer, social welfare can thereby be reduced. However, it is only the razor-vs.-blade pricing scheme that may become unprofitable. Indeed, if the manufacturer makes positive margins on the platform, then the availability of user-developed add-ons can have a positive effect: it can increase the value of the platform to users, and so allow manufacturers to charge higher margins on it and/or sell more units. Jeppesen (2004) finds that this is in fact the outcome when users introduce free game modifications (called mods) operating on proprietary game software platform products (called engines) sold by game manufacturers. Even though the game manufacturers also sell mods commercially that compete with free user mods, many provide active support for the development and diffusion of user mods built on their proprietary game engines, because they find that the net result is increased sales and profits.

Public Policy Choices

If innovation by users is welfare enhancing and is also significant in amount and value, then it makes sense to consider the effects of public policy on user innovation. An important first step would be to collect better data. Currently, much innovation by users—which may in aggregate turn out to be a very large fraction of total economic investment in innovation—goes uncounted or undercounted. Thus, innovation effort that is volunteered by users, as is the case with many contributions to open source software, is currently not recorded by governmental statistical offices. This is also the case for user innovation that is integrated with product and service production. For example, much process innovation by manufacturers occurs on the factory floor as they produce goods and simultaneously learn

how to improve their production processes. Similarly, many important innovations developed by surgeons are woven into learning by doing as they deliver services to patients.

Next, it will be important to review innovation-related public policies to identify and correct biases with respect to sources of innovation. On a level playing field, users will become a steadily more important source of innovation, and will increasingly substitute for or complement manufacturers' innovation-related activities. Transitions required of policy making to support this ongoing evolution are important but far from painless. To illustrate, we next review issues related to the protection intellectual property, related to policies restricting product modifications, related to source-biased subsidies for R&D, and related to control over innovation diffusion channels.

Intellectual Property

Earlier, when we explored why users might freely reveal their innovations, we concluded that it was often their best *practical* choice in view of how intellectual property law actually functions (or, often, does not function) to protect innovations today. For example, recall from chapter 6 that most innovators do not judge patents to be very effective, and that the availability of patent grant protection does not appear to increase innovation investments in most fields. Recall also that patent protection is costly to obtain, and thus of little value to developers of minor innovations—with most innovations being minor. We also saw that in practice it was often difficult for innovators to protect their innovations via trade secrecy: it is hard to keep a secret when many others know similar things, and when some of these information holders will lose little or nothing from freely revealing what they know.

These findings show that the characteristics of present-day intellectual property regimes as actually experienced by innovators are far from the expectations of theorists and policy makers. The fundamental reason that societies elect to grant intellectual property rights to innovators is to increase private investment in innovation. At the same time, economists have long known that there will be social welfare losses associated with these grants: owners of intellectual property will generally restrict the use of their legally protected information in order to increase private profits. In other words, intellectual property rights are thought to be good for innova-

tion and bad for competition. The consensus view has long been that the good outweighs the bad, but Foray (2004) explains that this consensus is now breaking down. Some—not all—are beginning to think that intellectual property rights are bad for innovation too in many cases.

The need to grant private intellectual property rights to achieve socially desirable levels of innovation is being questioned in the light of apparent counterexamples. Thus, as we saw earlier, open source software communities do not allow contributing innovators to use their intellectual property rights to control the use of their code. Instead, contributors use their authors' copyright to assign their code to a common pool to which all—contributors and non-contributors alike—are granted equal access. Despite this regime, innovation seems to be flourishing. Why? As we saw in our earlier discussions of why innovators might freely reveal their innovations, researchers now understand that significant private rewards to innovation can exist independent of intellectual property rights grants. As a general principle, intellectual property rights grants should not be offered if and when developers would seek protection but would innovate without it.

The debate rages. Gallini and Scotchmer (2002) assert that “intellectual property is the foundation of the modern information economy” and that “it fuels the software, lifesciences and computer industries, and pervades most other products we consume.” They also conclude that the positive or negative effect of intellectual property rights on innovation depends centrally on “the ease with which innovators can enter into agreements for rearranging and exercising those rights.” This is precisely the rub from the point of view of those who urge that present intellectual property regimes be reconsidered: it is becoming increasingly clear that in practice rearranging and exercising intellectual property rights is often difficult rather than easy. It is also becoming clear that the protections afforded by existing intellectual property law can be strategically deployed to achieve private advantage at the expense of general innovative progress (Foray 2004).

Consider an effect first pointed out by Merges and Nelson (1990) and further explored as the “tragedy of the anticommons” by Heller (1998) and Heller and Eisenberg (1998). A resource such as innovation-related information is prone to underuse—a tragedy of the anticommons—when multiple owners each have a right to exclude others and no one has an effective privilege of use. The nature of the patent grant can lead to precisely this type of

situation. Patent law is so arranged that an owner of a patent is not granted the right to practice its invention—it is only granted the right to exclude others from practicing it. For example, suppose you invent and patent the chair. I then follow by inventing and patenting the rocking chair—implemented by building rockers onto a chair covered by your patent. In this situation I cannot manufacture a rocking chair without getting a license from you for the use of your chair patent, and you cannot build rocking chairs either without a license to my rocker patent. If we cannot agree on licensing terms, no one will have the right to build rocking chairs.

In theory and in a world of costless transactions, people could avoid tragedies of the anticommons by licensing or trading their intellectual property rights. In practice the situation can be very different. Heller and Eisenberg point specifically to the field of biomedical research, and argue that conditions for anticommons effects do exist there. In that field, patents are routinely allowed on small but important elements of larger research problems, and upstream research is increasingly likely to be private. “Each upstream patent,” Heller and Eisenberg note, “allows its owner to set up another tollbooth on the road to product development, adding to the cost and slowing the pace of downstream biomedical innovation.”

A second type of strategic behavior based on patent rights involves investing in large portfolios of patents to create “patent thickets”—dense networks of patent claims across a wide field (Merges and Nelson 1990; Hall and Ham Ziedonis 2001; Shapiro 2001; Bessen 2003). Patent thickets create plausible grounds for patent infringement suits across a wide field. Owners of patent thickets can use the threat of such suits to discourage others from investing research dollars in areas of technical advance relevant to their products. Note that this use of patents is precisely opposite to policy makers’ intentions to stimulate innovation by providing ways for innovators to assert intellectual property rights. Indeed, Bessen and Hunt (2004) have found in the field of software that, on average, as firm’s investments in patent protection go up, their investments in research and development actually go down. If this relationship proves causal, there is a reasonable explanation from the viewpoint of private profit: corporations that can use a patent thicket to deter others’ research in a field might well decide that there is less need to do research of their own.

Similar innovation-retarding strategies can be applied by owners of large collections of copyrighted work in the movie, publishing, and software

fields. Copyright owners can prevent others from building new works on characters (e.g. Mickey Mouse) that are already familiar to customers. The result is that owners of large portfolios of copyrighted work can gain an advantage over those with no or small portfolios in the creation of derivative works. Indeed, Benkler (2002) argues that institutional changes strengthening intellectual property protection tend to foster concentration of information production in general. Lessig (2001) and Boldrin and Levine (2002) arrive at a similarly negative valuation of overly strong and lengthy copyright protection.

These types of innovation-discouraging effects can affect innovation by users especially strongly. The distributed innovation system we have documented consists of users each of whom might have only a few innovations and a small amount of intellectual property. Such innovators are clearly hurt differentially by a system that gives advantage to the owners of large shares of the intellectual property in a field.

What can be done? A solution approach open to policy makers is to change intellectual property law so as to level the playing field. But owners of large amounts of intellectual property protected under the present system are often politically powerful, so this type of solution will be difficult to achieve.

Fortunately, an alternative solution approach may be available to innovators themselves. Suppose that many elect to contribute the intellectual property they individually develop to a commons in a particular field. If the commons then grows to contain reasonable substitutes for much of the proprietary intellectual property relevant to the field, the relative advantage accruing to large holders of this information will diminish and perhaps even disappear. At the same time and for the same reason, the barriers that privately held stocks of intellectual property currently may raise to further intellectual advance will also diminish. Lessig supports this possibility with his creation and publication of standard "Creative Commons" licenses on the website creativecommons.org. Authors interested in contributing their work to the commons, perhaps with some restrictions, can easily find and adopt an appropriate license at that site.

Reaching agreement on conditions for the formation of an intellectual commons can be difficult. Maurer (2005) makes this clear in his cautionary tale of the struggle and eventual failure to create a commons for data on human mutations. However, success is possible. For example, an extensive

intellectual commons of software code is contained and maintained in the many open source software projects that now exist.

Interesting examples also exist regarding on the impact a commons can have on the value of intellectual property innovators seek to hold apart from it. Weber (2004) recounts the following anecdote: In 1988, Linux developers were building new graphical interfaces for their open source software. One of the most promising of these, KDE, was offered under the General Public License. However, Matthias Ettrich, its developer, had built KDE using a proprietary graphical library called Qt. He felt at the time that this could be an acceptable solution because Qt was of good quality and Troll Tech, owner of Qt, licensed Qt at no charge under some circumstances. However, Troll Tech did require a developer's fee be paid under other circumstances, and some Linux developers were concerned about having code not licensed under the GPL as part of their code. They tried to convince Troll Tech to change the Qt license so that it would be under the GPL when used in free software. But Troll Tech, as was fully within its rights, refused to do this. Linux developers then, as was fully within their rights, began to develop open source alternatives to Qt that could be licensed under the GPL. As those projects moved toward success, Troll Tech recognized that Qt might be surpassed and effectively shut out of the Linux market. In 2000 the company therefore decided to license Qt under the GPL.

Similar actions can keep conditions for free access to materials held within a commons from degrading and being lost over time. Chris Hanson, a Principal Research Scientist at MIT, illustrates this with an anecdote regarding an open source software component called ipfilter. The author of ipfilter attempted to "lock" the program by changing licensing terms of his program to disallow the distribution of modified versions. His reasoning was that Ipfilter, a network-security filter, must be as bug-free as possible, and that this could best be ensured by his controlling access. His actions ignited a flame war in which the author was generally argued to be selfish and overreaching. His program, then an essential piece of BSD operating systems, was replaced by newly written code in some systems within the year. The author, Hanson notes, has since changed his licensing terms back to a standard BSD-style (unrestricted) license.

We will learn over time whether and how widely the practice of creating and defending intellectual commons diffuses across fields. There obviously can be cases where it will continue to make sense for innovators, and for

society as well, to protect innovations as private intellectual property. However, it is likely that many user innovations are kept private not so much out of rational motives as because of a general, not-thought-through attitude that “we do not give away our intellectual property,” or because the administrative cost of revealing is assumed to be higher than the benefits. Firms and society can benefit by rethinking the benefits of free revealing and (re)developing policies regarding what is best kept private and what is best freely revealed.

Constraints on Product Modification

Users often develop prototypes of new products by buying existing commercial products and modifying them. Current efforts by manufacturers to build technologies into the products they sell that restrict the way these products are used can undercut users’ traditional freedom to modify what they purchase. This in turn can raise the costs of innovation development by users and so lessen the amount of user innovation that is done. For example, makers of ink-jet printers often follow a razor-and-blade strategy, selling printers at low margins and the ink cartridges used in them at high margins. To preserve this strategy, printer manufacturers want to prevent users from refilling ink cartridges with low-cost ink and using them again. Accordingly, they may add technical modifications to their cartridges to prevent them from functioning if users have refilled them. This manufacturer strategy can potentially cut off both refilling by the economically minded and modifications by user-innovators that might involve refilling (Varian 2002). Some users, for example, have refilled cartridges with special inks not sold by printer manufacturers in order to adapt ink-jet printing to the printing of very high-quality photographs. Others have refilled cartridges with food colorings instead of inks in order to develop techniques for printing images on cakes. Each of these applications might have been retarded or prevented by technical measures against cartridge refilling.

The Digital Millennium Copyright Act, a legislative initiative intended to prevent product copying, may negatively affect users’ abilities to change and improve the products they own. Specifically, the DMCA makes it a crime to circumvent anti-piracy measures built into most commercial software. It also outlaws the manufacture, sale, or distribution of code-cracking devices used to illegally copy software. Unfortunately, code cracking is also

a needed step for modification of commercial software products by user-innovators. Policy makers should be aware of “collateral damage” that may be inflicted on user innovation by legislation aimed at other targets, as is likely in this case.

Control over Distribution Channels

Users that innovate and wish to freely diffuse innovation-related information are able to do so cheaply in large part because of steady advances in Internet distribution capabilities. Controls placed on such infrastructural factors can threaten and maybe even totally disable distributed innovation systems such as the user innovation systems documented in this book. For example, information products developed by users are commonly distributed over the Internet by peer-to-peer sharing networks. A firm that owns both a channel and content (e.g., a cable network) may have a strong incentive to shut out or discriminate against content developed by users or others in favor of its own content. The transition from the chaotic, fertile early days of radio in the United States when many voices were heard, to an era in which the spectrum was dominated by a few major networks—a transition pushed by major firms and enforced by governmental policy making—provides a sobering example of what could happen (Lessig 2001). It will be important for policy makers to be aware of this kind of incentive problem and address it—in this case perhaps by mandating that ownership of content and ownership of channel be separated, as has long been the case for other types of common carriers.

R&D Subsidies and Tax Credits

In many countries, manufacturing firms are rewarded for their innovative activity by R&D subsidies and tax credits. Such measures can make economic sense if average social returns to innovation are significantly higher than average private returns, as has been found by Mansfield et al. (1977) and others. However, important innovative activities carried out by users are often not similarly rewarded, because they tend to not be documentable as formal R&D activities. As we have seen, users tend to develop innovations in the course of “doing” in their normal use environments. Bresnahan and Greenstein (1996a) make a similar point. They investigate the role of “co-invention” in the move by users from mainframe to client-server architecture.¹ By “co-invention” Bresnahan and Greenstein mean organizational

changes and innovations developed and implemented by users that are required to take full advantage of a new invention. They point out the high importance that co-invention has for realizing social returns from innovation. They consider the federal government's support for creating "national information infrastructures" insufficient or misallocated, since they view co-invention as the bottleneck for social returns and likely the highest value locus for invention.

Efforts to level the playing field for user innovation and manufacturer innovation could, of course, also go in the direction of lessening R&D subsidies or tax credits for all rather than attempting to increase user-innovators' access to subsidies. However, if directing subsidies to user-innovators seems desirable, social welfare will be best served if policy makers link them to free revealing by user-innovators as well as or instead of tying them to users' private investments in the development of products for exclusive in-house use. Otherwise, duplication of effort by users interested in the same innovation will reduce potential welfare gains.

In sum, the welfare-enhancing effects found for freely revealed user innovations suggest that policy makers should consider conditions required for user innovation when creating policy and legislation. Leveling the playing field for user-innovators and manufacturer-innovators will doubtless force more rapid change onto manufacturers. However, as will be seen in the next chapter, manufacturers can adapt to a world in which user innovation is at center stage.

