Appendix 2 | Modeling Free Innovation's Impacts on Markets and Welfare

In chapter 6, I summarized and discussed the findings of the modeling presented in Gambardella, Raasch, and von Hippel (2016). The model itself is significantly richer than a non-mathematical summary can convey, and so in this appendix I reproduce the original version of our model "set-up" information, the mathematical version of the model itself, and related findings exactly as presented in sections 4 and 5 and the appendixes of Gambardella, Raasch, and von Hippel (2016). Before reading this appendix, readers might wish to review chapter 6 above for contextual information not repeated below.

In my description of the model and the findings in chapter 6, I changed the term "user innovators" (used in Gambardella, Raasch, and von Hippel 2016) to "free innovators." However, I have kept the original term, "user innovator," in this appendix. I have done so because in the research article we defined user innovators as having exactly the same range of possible self-reward types as free innovators. The only difference is that in the article it was assumed that user innovators always reaped *some* level of self-reward from personal use, perhaps in addition to other possible types of self-rewards. This assumed pattern is likely to be very generally the case among real-world free innovators. Recall that some level of use motivation was present in all the free innovator cluster data shown in figure 2.1 of the present book.

Editor's note As was explained above, the remainder of this appendix is quoted from Gambardella, Raasch, and von Hippel (2016). The material is reproduced here as supplied by the author. The original numbering of sections and subsections has been retained. The figure, which was numbered 2 in the original paper, is referred to here simply as "the figure."

Section 4. Model set-up and findings

4.1 User types and 'tinkering surplus'

We divide a producer's potential market into two types of users: *innovating users* and *non-innovating users*. *Innovating users* find it viable to develop and self-provision innovative designs related to the producer product, e.g., improvements, customizations, and complements. They can also viably self-provision home-made copies of the producer product itself, and so can choose whether to buy the product from a firm or to make it themselves. *Non-innovating users* do not have a viable option of innovating. Their costs may be too high, for example, because they lack needed skills or access to tools, or because they have a high opportunity cost for their time. However, it is viable for non-innovating users to make copies and self-provision products based on designs developed by user innovators at some level of quality ranging from equal to innovating users down to zero.

The share of innovating users is σ, and we regard this share as exogenous and static; users cannot change their type. For simplicity, we normalize the size of the market to 1, so that σ and $1 - \sigma$ are also the number of users of each respective type.

With respect to the utility users derive from innovating, we note that empirical research finds that innovating users derive utility *both* from using the innovation they have created, and from innovation process benefits they gain from engaging in the innovation process itself, such as fun and learning (Lakhani and Wolf 2005; Franke and Schreier 2010; Raasch and von Hippel 2013). Users seek to maximize their utility from innovating, which we call h, by determining the optimal amount of resources, such as time t, to devote to innovation projects,

$$Max(t)h \equiv \chi + (\phi^{1-\alpha}/\alpha)x^{1-\alpha}t^{\alpha} + 1 - t. \quad (1)$$

In equation (1), the parameter χ represents a user innovator's utility, net of all innovation-related costs, from go-it-alone innovation projects, i.e., when producers do nothing to support him. The second term of (1) represents the user innovator's additional utility when a firm conducts x projects to support his endeavors. Examples of such support are the development of design tools for the use of users, and gamification to make product design activities more enjoyable to users. The parameter $\alpha \in (0,1)$ captures whether innovating users' utility is mostly determined by the time they invest (high α) or by the extent of firm support (low α). The parameter $\phi > 0$ captures the productivity of this process. The last term, $1 - t$, captures the value of the user innovator's remaining time that he can spend on other matters, when the total time he has available is normalized to 1 and he has decided to spend t on innovation projects.

We derive from (1) that the user's utility-maximizing time investment in innovation is $t = \phi x$, which yields utility $h = \chi + (1 - \alpha/\alpha)\phi x + 1$. We call this expression capturing users' net benefit from innovating the *tinkering surplus* (TS), where TS is the aggregate net benefit that all users gain from innovating and self-provisioning. It consists of benefits from the use of the self-provisioned innovation, plus innovation process benefits, as mentioned above, minus costs. When the investment of firms in user innovation support is zero, innovating users still get their go-it-alone tinkering surplus, $h = \chi + 1 > 0$. If firms do invest ($x > 0$), TS increases as a function of the level of that investment.

4.2 Shared vs. producer-only innovation

We decompose the value that all buyers derive from the producer product into two parts: value v that they derive from features and components that only the producer firm will develop and produce, and value b that buyers derive from features and components that can be developed and produced by firms *and* users, jointly or in isolation.

Features that only producers will find viable to develop include those that offer limited value to many individual users. No individual user would find it viable to develop such a feature, but producers can aggregate demand across buyers and thereby recoup their investment (Baldwin and von Hippel 2011). Features in this category may include, e.g., product engineering for greater durability and ease of use, a more elaborate design, a manual to accompany the product, etc. In contrast, features b that both individual users (typically "lead users") and producers can viably develop require smaller investments, compensated for by larger benefits to individual user innovators. They provide high functional novelty and solve important, hitherto unmet user needs (von Hippel 2005). As the needs of lead users foreshadow demand in the market at large (cf. definition of lead user), non-innovating users too will predictably benefit from solutions to these problems with the passage of time.

We assume that all users tend to have more similar assessments of the features we call b, that innovative users may get involved in developing, than of the features v that the producer has to develop on its own. Capturing this idea of less heterogeneity with regard to b but simplifying our analysis, we assume that users differ only in their valuations of v ($v \sim U[0, 1]$) whereas they all like b to the same degree. In our model of innovation and production by users and producers, we focus on innovations of type b, following our assumption that producers are the only ones to invest in v. Innovations with regard to b are assumed to depend on two activities.

First, the volume of innovations of type b depends on the aggregate effort T exerted by all innovating users, to the extent that it is useful to the firm (e.g., net of redundancy). To streamline our analysis, we assume that the aggregate usable effort is simply proportional to the total efforts t of the σ innovating users, that is $T = \gamma'\sigma\, t$, $\gamma' > 0$. (We could use more complex aggregations, allowing for increasing or diminishing returns to the number of innovating users, but our results would remain materially unchanged.) Assuming identical innovating users, and employing the optimal expression for t, $t = \phi x$, we obtain aggregate user effort

$$T = \gamma\sigma x$$

where $\gamma = \gamma'\phi$ comprises any factor that raises the ability of the firms to take advantage of the productivity of the innovating users' efforts to improve b. As explained earlier, the firm can influence aggregate user effort T through x projects to develop tools and platforms that support and leverage innovating users. The projects affect the time t users want to spend on innovation projects, which then affects the value of the innovative product b via aggregate effort T.

Second, innovations of type b are a function of some commitment of resources Y carried out by the firm. To fix ideas, Y can be commercial

R&D projects or any other product creation or development activity. We define

$$Y = \xi(1 - s)y, \quad \xi \geq 0$$

where y is the total number of innovation projects of the firm. The firm allocates a share s to projects that support innovating users, that is $x = sy$, and the remainder, $(1 - s)y$, goes to traditional commercial R&D projects (either in-house or external). Projects that support innovating users are of little commercial value, per se, but indirectly produce value by attracting more user innovation activities. The parameter ξ measures the productivity of the firm's commercial R&D.

Taking into account these two drivers of innovation—aggregate user effort T and producer R&D activity Y—let the value of the innovative product to users be

$$b = (T^{\beta} + Y^{\beta})^{1/\beta}, \qquad \beta > 0,$$

which we can rewrite as

$$b = \left[\tau^{\beta} s^{\beta} + \xi^{\beta}(1-s)^{\beta}\right]^{1/\beta} y = \tilde{b}y$$

where $\tau \equiv \gamma\,\sigma$ and $\tilde{b}\left[\tau^{\beta} s^{\beta} + \xi^{\beta}(1-s)^{\beta}\right]^{1/\beta}$ is the productivity of all the firm's y projects taken together.

4.3 User and producer innovation activity as substitutes or complements

The parameter β plays an important role in our analysis. It captures two options that firms can choose from, each of which involves a distinct form of organizing tasks and resources for innovation. The first option is such that the efforts of innovating users, T, and those of the producer, Y, are substitutes. Take, for instance, the writing of new software code. Suppose that both the producer and users can work on each of two tasks, (1) novel functionality, and (2) the creation of convenience-enhancing features such as "user friendly" installation scripts. The more effort the producer spends on each of these tasks, the lower the innovation impact that users can make, and vice versa. One effort tends to substitute for the other. In our model, this situation is captured by $\beta > 1$, which implies that the marginal impact of T on b decreases as Y increases and vice versa.

The second option, in contrast, structures R&D for complementarity between user and producer innovation activities. In our example, suppose that users write novel code and producers develop "convenience features." The more effort users put into coding, the higher the impact that producers can make, and vice versa. In our model, this situation is described by $0 < \beta < 1$, which implies that the marginal impact of T on b increases as Y increases and vice versa. Research has shown that user innovators tend to focus on developing innovations providing novel functionality, and producers on developing innovations that increase product reliability and user convenience (Riggs and von Hippel 1994; Ogawa 1998). A good example in the software field is RedHat. That firm's commercial offerings are based on open source software code such as Linux and Apache software, developed by users, to which RedHat adds convenience features such as "easy installation" software scripts.

To streamline our analysis, we assume that each firm can pick its preferred innovation option, but not the specific level of β. A fully endogenous β would add complexity without substantial new insight. In practice, its value will depend on the industry in question, the technologies available to the firm, and best practices for integrating innovating users in R&D.

4.4 Individual market demands of innovating users and non-innovating users

Next, we need to understand the demand for the producer product from non-innovating users and from innovating users given user contestability, user-created complements and spillovers, i.e., the different types of interactions that we developed in section 3 [of this paper].

Starting with innovating users, we expect that they will buy the product from a firm only if their consumer surplus is positive and exceeds their surplus from self-provisioning, i.e., if

$$v + b - p + h \geq \lambda b + h, \; v \sim \mathrm{U}[0, 1], \qquad 0 \leq \lambda \leq 1. \tag{2}$$

The term $v + b - p$ is the consumer surplus, where $v + b$ is our value decomposition of the producer product (cf. section 4.2) and p is its price. In case of self-provisioning, a user innovator will not get utility v, which is provided by the firm only. Of utility b that all innovating users co-create with the firm, he will get only the "walk-away value" λb that he can realize by learning from this co-creation process and trying to build features akin to b on his own. The quality $0 \leq \lambda \leq 1$ of his self-provisioned version of b will depend on several factors, such as the extent and format of information spillovers from the firm to the user innovator, his "absorptive capacity" for the spillovers, and his skills to build the information into a usable artifact. In the case of software programming, for instance, where the producer opens up his source code for users to co-develop, λ will be close to 1, if and as the essential design information required to replicate functionality b is fully revealed. In this example, if the producer shares only part of his source code, λ is depressed accordingly.

Finally, recall the user innovator's surplus h from her own innovation activities, including those extensions and customizations that the firm is not interested in. The user innovator is assumed to get this surplus h—the tinkering surplus—regardless of whether she buys the producer product or not.

Recall that non-innovating users simply buy a producer-provisioned product via the market or, to the extent that they are able, can elect to replicate a design developed and then shared peer to peer by a user innovator. Building on what we said earlier about v, b and p as constituents of demand, we expect that non-innovating users will buy on the market if

$$v + b - p + \mu' h \geq \mu b + \mu' h, \qquad 0 \leq \mu, \qquad \mu' \leq 1 \tag{3}$$

and self-provision otherwise.

The parameters μ and μ' in equation (3) capture the non-innovating users' ability to obtain knowledge of the innovating users' designs (which will depend on the innovating users' propensity to diffuse design information), to replicate them, and to benefit from them. While μ' refers to a non-innovating user's ability to benefit from an individual user innovator whose design they adopt, μ captures their trickle-down benefits from what the user innovator has learned from the producer as well as other innovating users during the co-creation process of b. Of course, when the non-innovating users buy from the firms they enjoy b incorporated in the firms' product, while they enjoy μb when they obtain the product from the innovating users through peer-to-peer diffusion. We expect non-innovating users to have imperfect knowledge of the innovating users' designs, to be less skilled at self-provisioning them and/or to benefit less from using them ($\mu \leq \lambda$ and $\mu' \leq 1$). With respect to imperfect knowledge and higher costs of self-provisioning, consider that innovating users may well regard careful design documentation for the benefit of potential adopters to be an unprofitable chore in the case of freely revealed designs (de Jong et al. 2015; von Hippel, DeMonaco, and de Jong 2016). With respect to lower levels of benefit, consider that the designs were developed to precisely suit the innovating users' individual tastes.

Finally, it is crucial to note the trade-off that our model implies for the producer: Firms benefit from learning from innovating users about how to make a better product *b* for both innovating and non-innovating customers; to that end they want to invest in *x* to involve users more extensively. At the same time, this comes at the cost of facilitating self-provisioning by both innovating and non-innovating users. As the producer invests in tools and toolkits, modularizes the product or reveals design knowledge such as source code to facilitate user innovation, he also makes it easier for both innovating and non-innovating users to self-provision rather than buy. Our model assumes that the producer cannot entirely avoid this side effect of enhanced user contestability, even while choosing a mode of supporting user innovation that best serves his goals.

4.5 Profit maximization by firms

The aggregate demanded quantity of $(1 - \sigma)$ non-innovating users and σ innovating users is

$$q = (1 - \sigma)(1 - p + (1 - \mu)b) + \sigma(1 - p + (1 - \lambda)b) = 1 - p + \eta b, \quad (4)$$

with

$$\eta \equiv (1 - \mu)(1 - \sigma) + (1 - \lambda)\sigma.$$

Solving for *p*, inverse demand is

$$p = 1 + \eta b - q. \quad (5)$$

With N symmetric firms in the market, aggregate demand is $q = \sum_{j=1}^{N} q_j$ and q/N is the demand faced by one firm. Firm profits Π_i are given by the number of units sold by firm i, q_i, times the profit margin, given by price p minus marginal cost of production φ, and minus the cost of y innovation projects,

$$\Pi_i = (p - \varphi)q_i - \kappa y^2, \qquad \kappa > 0 \tag{6}$$

where we assume diminishing returns to running y projects.

In order to maximize profits, firms make several interrelated decisions in the following sequence: First, they decide on the organization of their R&D. Specifically, they pick one of two options available to them: the organization of R&D such that user and producer inputs, T and Y, are substitutes ($\beta > 1$) or the organization for complementarity ($0 < \beta < 1$). It will take firms longer to change their organizational structure and capabilities in R&D than to change the number of projects, which is why we model this is as the first choice. Next, the firms pick their total number of R&D-related projects (y). Then they decide on the share of projects ($1 - s$) to allocate to traditional producer R&D. The remainder of the projects, share s, will be devoted to user-innovation support and thus indirectly increase the flow of new product ideas available to the firm. Finally, firms decide on the quantity to produce and sell on the market (q_i).

We use backward induction to derive the producers' optimal decisions. In this section, we look at the optimal choices of q_i, s, and y, in this order. In section 4.7, we will study the choice of innovation mode (β).

Choice of q_i. We take the derivative of (6) with regard to output quantity (q_i) and obtain *foc*: $1 + \eta b - \varphi - \sum_{j=1}^{N} q_j - q_i = 0$. In symmetric equilibrium, this produces profit-maximizing quantity, price and profits

$$q_i = (1 + \eta b - \varphi)/(N + 1) \tag{7a}$$

$$p = (1 + \eta b - \varphi)/(N + 1) + \varphi \tag{7b}$$

$$\Pi_i = (p - \varphi)^2 - \kappa y^2 = [(1 + \eta b - \varphi)/(N + 1)]^2 - \kappa y^2. \tag{7c}$$

Choice of *s*. To determine the share of the firms' projects aimed at supporting user innovation, *s*, we maximize $\tilde{b}^U = (\xi^\theta + \tau^\theta)^{1/\theta}$ yielding *foc* $\tau^\beta \beta s^{\beta-1} - \beta\xi^\beta(1-s)^{\beta-1} = 0$. To determine the optimal *s*, a case distinction is required. In the case of complementarity between user efforts and producer R&D, i.e., if $0 < \beta < 1$, the *soc* is negative, which implies that there is an intermediate project allocation $0 < s < 1$ to user support that maximizes innovation output $\tilde{b}$ (specifically $s = \tau^\theta/(\xi^\theta + \tau^\theta)$, with $\theta \equiv \beta/(1-\beta)$). As can be seen from the expression for τ, this optimal project share allocated to user innovation supports increases in the share of innovating users in the market and their productivity in terms of commercially valuable ideas ($s_\sigma, s_\gamma > 0$, where from now on we use subscripts to denote derivatives), and decreases with the productivity of producer R&D ($s_\xi < 0$). In the case of substitution between user and producer innovation efforts, i.e., if $\beta > 1$, the *soc* is positive, which implies that the optimal allocation to user support, *s*, is either 0 or 1, depending on whether the productivity of the user contribution in $\tilde{b}$, that is τ, is greater or smaller than the productivity of the firm contribution, ξ.

Choice of *y*. The *foc* of (7c) with respect to *y* is $2(1+\eta b-\varphi)\eta\tilde{b}/(N+1)^2 - 2\kappa y = 0$, which yields $y = (1-\varphi)z/[\kappa(N+1)^2 - z^2)$, where $z \equiv \eta\tilde{b}$. Note that the *soc* implies $\kappa(N+1)^2 - z^2 > 0$ such that the profit-maximizing investment *y* is always positive. It is also easy to see that *y* increases with *z*. κ

4.6 The producer vs. user-augmented innovation modes

From our findings from the previous section relating to the distribution of innovation projects by the firm (*s*), we see that there are two modes of innovating, and that firms will want to choose between them. The first mode is characterized by $\beta > 1$ and $s = 0$. That is, in this mode firms choose to organize their R&D such that user and producer efforts are substitute inputs and then allocate their entire budget to their own commercial R&D efforts, not supporting user innovation activities in any way. We call this the producer (P) *innovation mode*. In this mode, firms ignore the innovating users and organize the creation of *b* solely around closed commercial R&D.

As a consequence of being closed, firms need not fear information spillovers to innovating users ($\lambda = 0$) and on to non-innovating users (μ

= 0). In the producer mode, therefore, the demands of the non-innovating users and the innovating users simplify to

$$v - p + b + \mu' h \geq \mu' h \tag{8}$$

$$v - p + b + h \geq h, \tag{8'}$$

respectively. At the same time, aggregate demand is (4), with $\eta = 1$ rather than $(1 - \mu)(1 - \sigma) + (1 - \lambda)\sigma$.

The second innovation mode is characterized by the firm organizing its R&D for complementarity with user innovators ($0 < \beta < 1$) and then making a positive investment in user innovation support (optimal $s = \tau^\theta/(\xi^\theta + \tau^\theta) > 0$). We call this the *user-augmented* (U) *mode*. In this mode, firms actively leverage user-created spillovers for innovation and organize their R&D to exploit the complementarity between the two sources of innovation. Users contribute to raising the use value b of the product, which enhances the demand of both the non-innovating users and the innovating users. At the same time, firms' support of innovating users creates user contestability with regard to features b ($\lambda, \mu \geq 0$).

To summarize, the trade-off between the U- vs. P-modes pivots on producers investing to facilitate user innovation and reap spillovers, but by this action simultaneously and unavoidably boosting user self-provisioning to a degree that may be small or large.

4.7 Choice of innovation mode (β)

Continuing our earlier process of backward induction to understand outcomes in markets with innovating users, we now consider the very first producer decision, the choice of innovation mode. Our goal is to understand under what conditions a producer will prefer the producer mode over the user-augmented mode, or *vice versa*. Additionally and importantly, we examine under what conditions the increasing prevalence of innovating users that we observe in many markets (cf. Baldwin and von Hippel 2011) renders user integration the profit-maximizing innovation strategy for producers.

Our first theorem below explains the choice of innovation mode by a producer firm. It establishes that, subject to two conditions, firms in markets with an increasing share of innovating users will find it in their own best interest to switch to the user-augmented mode. In switching, firms are aware that they are strengthening user contestability, but also realize that, overall, this is more profitable than a closed innovation approach.

To find the profit-maximizing mode of innovation, it is convenient to rewrite expression (7c) for the profits of the firms as

$$\Pi = [(1 + zy - \varphi)/(N + 1)]^2 - \kappa y^2. \quad (9)$$

This expression captures profits in both the P- and U-modes, which differ only in z. (In particular, in the P-mode $z^P = \eta^P \tilde{b}^P$, with $\eta^P = 1$ and $\tilde{b}^P = \xi$; in the U-mode $z^U = \eta^U \tilde{b}^U$, with $\eta^U = (1 - \mu)(1 - \sigma) + (1 - \lambda)\sigma$ and $\tilde{b}^U = (\xi^\theta + \tau^\theta)^{1/\theta}$.) Given the optimal choices of s and y, as derived in section 4.5, this implies that $\Pi^P \geq \Pi^U$ if and only if $z^P \geq z^U$. In other words, we can check whether profits are higher in the P- or U-mode simply by checking whether z is higher in one or the other.

We find that, when there are very few innovating users (σ close to zero), profits in the P-mode are always higher than profits in the U-mode ($\Pi^P > \Pi^U$). Thus, when there are very few innovating users, firms choose the P-mode. The intuition is that, from the firms' perspective, the user innovation spillovers that they can harvest, the upside of conducting projects x to support user innovation, are low. At the same time, the downside is considerable, as the information and tools that the firms supply to the few innovating users can enable innovating users to develop a competing design and share it peer to peer, knocking off a good part of the producer's demand. The magnitude of this loss, and thus the downside of switching to the U-mode, will depend on λ and μ, users' ability to self-provision b.

As the share of innovating users increases, profits stay the same in the P-mode but increase in the U-mode. (This is true under two conditions that we will explain below.) Firms will switch from the P- to the U-mode when the share of innovating users is larger than a threshold σ^*, beyond which $\Pi^U > \Pi^P$. This is illustrated in [the figure].

The first condition relates to λ and μ. When user contestability is very weak (as indicated by the uppermost curve for which $\lambda = \mu = 0$), the producer can switch to the user-augmented mode free of risk. On this curve, when the share of innovating users is $\sigma = 0$, profits are equal for both modes of innovating. Then, as σ increases, the U-mode outpaces the P-mode in terms of firm profits. Intuitively, in this case firms

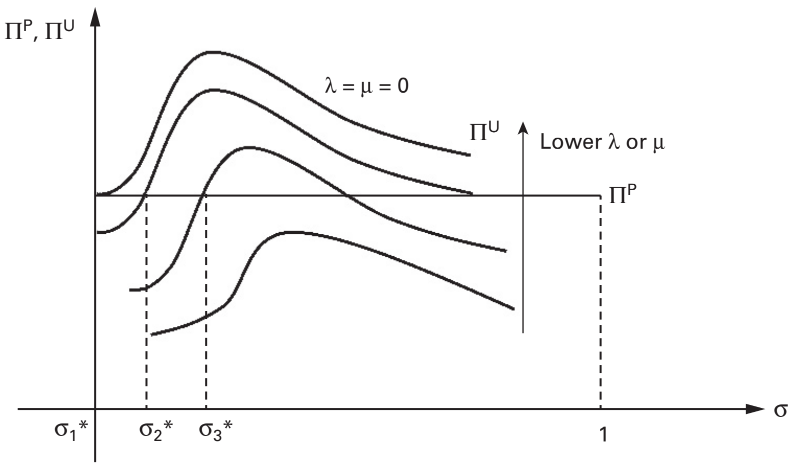

Firms' profits under the U- and P-modes.

benefit from the contribution of innovating users without risking the rise of self-provisioning and concomitant reduction of demand for the firms' product. When user contestability is more pronounced (as illustrated by the second and third curve), we see that the threshold σ* at which the switch to the U-mode can occur shifts to the right—that is, a higher share of innovating users in the market is needed for the producer to prefer the U-mode. If spillovers λ or μ are very large, as illustrated by the bottom curve, a switch to the U-mode will never be attractive to firms.

The second necessary condition for the switch to occur is that the complementarity between user and producer efforts *T* and *Y* must be strong enough. Specifically, $\theta < 1$ (i.e., $0 < \beta < ½$) must hold.[1] In other words, the contribution of the innovating users must be strong enough to trigger a significant increase in *b* that outweighs the negative impact on profits from intensified user contestability; otherwise firms will prefer to stay in the producer mode.

Theorem (Choice of mode). *If the innovative contribution of the innovating users is sizable ($0 < \beta < ½$) and user-contestability (λ and μ) is not too high, a critical mass of innovating users ($\sigma > \sigma^*$) makes profit-maximizing firms prefer the user-augmented mode of innovating to the producer mode.*

Proof. See Appendix A [of this paper].

It should be noted that while firms may find it profitable to switch to the U-mode at threshold σ^*, they may switch back again to the P-mode at a high σ. (As illustrated in [the figure], Π^U reaches a maximum and then declines, potentially even falling below the Π^P line.) This is particularly likely at higher levels of λ and μ (cf. [the figure]). The reason is the following: By our assumption of $\lambda \geq \mu$ innovating users are more capable than non-innovating users of self-provisioning, i.e., they exhibit a superior outside option and thus lower demand for the product of the firm. When the share of innovating users σ gets quite large, this not only means extensive user innovation spillovers to firms but also implies that the share of non-innovating users—those who benefit the most from these spillovers by getting to buy a superior product—is small. Having many innovating users implies having low demand, particularly if λ is large. This detracts from the attractiveness of the U-mode and may make firms prefer to switch back to the P-mode where they can better capture demand. We will leave this issue for future research to investigate in more detail, since our core objective is to understand the initial switch from the producer mode to the user-augmented mode when the prevalence of innovating users increases.

4.8 Welfare and policy

In this final section, we consider the welfare implications of firms choosing either to "go it alone" in the producer mode of innovating or to integrate user inputs in the user-augmented mode. We need to understand whether firms' choice of mode is efficient from a societal perspective and if not, whether policy is likely to improve economic outcomes.

Calculations of social welfare that include user innovation are different from the standard mode of calculating welfare. Conventionally, social welfare is calculated as profits (PS) plus consumer surplus (CS). When innovating users develop and build a new product for their own use, welfare calculations must be modified to include their full costs and benefits. In particular, we need to take into account their tinkering surplus TS, which is the aggregate net benefit that all users gain from self-provisioning, if they choose to do so. To give an example, if a user self-provisions a newly designed product at a cost of 10 dollars and receives a monetized use value of 30 dollars, her tinkering surplus equals 20 dollars. Recall that benefits to tinkering can also accrue in the form of process value (Franke and Schreier 2010; Raasch and von Hippel 2013), e.g., enjoyment of or learning from the innovation process itself, or social status in the user community. Our model is agnostic to the composition of these benefits. It only presumes them not to be profit based, in line with the definition of a user innovator. We will consider generalizations of this aspect in the discussion section.

Incorporating these considerations, then, welfare in markets containing both user and producer innovators should be computed as

$$W = \text{PS} + \text{CS} + \text{TS} \tag{10}$$

where PS and CS are the standard producer and consumer surplus and TS is the tinkering surplus. How significant is the omission of the tinkering surplus in conventional analyses? The answer depends on the extent of user self-provisioning in a market. If many users self-provision (as is common across an increasing range of markets, especially markets for digital products, cf. Baldwin and von Hippel 2011), the omission can be substantial. In some cases, it may dwarf traditional components of welfare.

In our model, the tinkering surplus for a user innovator equals h, while for non-innovating users it is $\mu' h$, which stems from their ability to tap into peer-to-peer diffusion from the innovating users. Computing the components of welfare as they accrue to producers (aggregate profits, PS), non-innovating users (consumer surplus, CS^{nui} plus tinkering surplus TS^{nui}) and innovating users (CS^{ui} plus tinkering surplus, TS^{ui}), we have

$$\text{PS} = N\Pi$$

$$\text{CS}^{\text{nui}} + \text{TS}^{\text{nui}} = (1 - \sigma)\ [(1 - p + (1 - \mu)b)^2/2 + \mu b + \mu' h]$$

$$\text{CS}^{\text{ui}} + \text{TS}^{\text{ui}} = \sigma\ [(1 - p + (1 - \lambda)b)^2/2 + \lambda b + h]$$

The first term is the aggregate profit of all the producers. The second term is the aggregate surplus of all non-innovating users, calculated from

$$(1-\sigma)\left[\int_{p-(1-\mu)b}^{1}(v-p+b+\mu' h)\,dv+\int_{0}^{p-(1-\mu)b}\mu b+\mu' h dv\right].$$

The third term derives from

$$\sigma\left[\int_{p-(1-\lambda)b}^{1}[v-p+b+h]dv+\int_{0}^{p-(1-\lambda)b}(\lambda b+h)dv\right].$$

These expressions will differ, depending on whether the U-mode or the P-mode is being chosen by firms.

Our analysis of welfare produces two main results that we summarize in two theorems. The first theorem states that, given our condition $0 < \beta < ½$, higher firm profits in the U-mode imply higher welfare in the U-mode, but the reverse is not true. That is, whenever firms' profits are higher in the U-mode, welfare is aligned; in contrast, when firms' profits are higher in the P-mode, welfare may not be aligned. Specifically, there are levels of σ, the share of innovating users in a market, such that profits are higher in the P-mode but welfare is higher in the U-mode. As a result, to the extent that the decision to switch belongs to the producers, as we modeled it, producers will remain in the P-mode even though the share of innovating users is substantial and social welfare would be better served in the U-mode. The reason is that firms do not internalize the key externalities of our model—that is, the increase in tinkering surplus (h) accruing to users because of firms' investment in user support (x) and also facilitation of self-provisioning that firms bestow on innovating users (λb) and, subsequently, non-innovating users (μb) even if they do not buy the product.

Theorem (Welfare)*. Under the conditions of the choice-of-mode theorem, if firms' profits are higher in the user-augmented mode, so is welfare, but the reverse is not true.*

Proof. See Appendix B [of this paper].

Our second result regards policy. We show that policies that increase the productivity of innovating users *can never reduce welfare*, provided that the costs of such policies do not outweigh their benefits. By contrast, policies that increase the productivity of R&D within firms *may reduce welfare*.

Examples of policies that raise the productivity of innovating users, γ, are subsidized access to design tools and maker-spaces. If innovating users become more productive, both profits and welfare rise under the U-mode but not under the P-mode, which, after all, does not leverage users' productivity. As firms' profits in the U-mode increase, they may come to exceed profits in the P-mode. We know from the previous theorem that if profits are higher in the U-mode, welfare is also higher.

Policies that increase the productivity of producer R&D, ξ, include R&D subsidies and tax exemptions as well as publicly funded applied R&D. Increases in firms' research productivity ξ raise profits in both the P- and the U-modes. We show that, unless complementarity between user and producer efforts is high, increases in ξ induce a larger increase in profits in the P-mode than in the U-mode.[2] This means that policies that support traditional producer R&D may induce a switch back to the P-mode. Since welfare is sometimes lower in the P-mode even while firms prefer it, increases in ξ may render the P-mode more attractive to the firms in spite of the fact that welfare is higher in the U-mode. In other words, such increases may induce a switch to the P-mode while welfare is higher in the U-mode, or prevent a welfare-increasing switch to the U-mode.

To summarize, policies that support producer innovation productivity ξ may reduce welfare. The mechanism is that such policies encourage firms to adopt a closed producer innovation mode while welfare may be higher in an open user-augmented mode. In contrast, policies that support the productivity of the innovating users can never reduce welfare. This is because they can only encourage a switch to the U-mode and this is never welfare reducing because whenever firms prefer the U-mode welfare is higher in this mode.

Theorem (Policy). *Under the conditions of the choice-of-mode theorem, policies that raise the productivity of innovating users, γ, encourage firms to adopt the user-augmented mode and can never reduce welfare. By contrast, if the complementarity between user and producer innovation activities,* T *and* Y*, is weak ($\beta > \beta^*$, with $\beta^* < ½$), policies that raise firms' research productivity, ξ, encourage firms to adopt the producer innovation mode, which may reduce welfare.*

Proof. See Appendix C [of this paper].

Section 5 and 5.1 Discussion

In this paper we analyzed the effects of user innovation by consumers on standard outcomes in markets for innovation. Our special focus was on understanding the implications of the increasing prevalence of innovating users (σ increasing from a low level), as found in many markets.

Our principal findings were three. First, as the share of innovating users in a market increases beyond a certain threshold, firms' profit-maximizing strategy is to switch from the traditional producer-only innovation approach to an innovation mode that harnesses user innovators. Subject to two intuitive conditions relating to the innovative and competitive impact of user activities, welfare is higher in this user-augmented mode than in traditional producer-only innovation mode. All of the constituencies—producers, innovating users, and non-innovating users—benefit.

Second, any firm that elects to switch to integrating innovating users definitely augments social welfare; but firms generally switch "too late." Thus, markets containing both user and producer innovators tend to fall short of their theoretical optimum in terms of value creation because producers are too slow, from a social welfare perspective, to embrace user innovation. Thus, producers' optimal R&D strategies yield a suboptimal division of innovative labor between users and producers at the societal level. Underlying this inefficiency there are externalities that the producer cannot capture, e.g., the "tinkering surplus" that accrues to users, a novel component of social welfare.

Third, policies that raise the productivity of innovating users encourage firms to switch to the user-augmented mode and can never reduce welfare. By contrast, policies that raise firms' research productivity encourage firms to switch back to the traditional producer-innovation mode and thereby may reduce welfare.

5.1 Assumptions, robustness and generalizability of findings

Our model rests on several assumptions that can be usefully investigated via further research.

First, as we mentioned at the start of the paper, innovating users are defined as individuals or firms developing innovations to use rather than sell. In this paper, we have focused on individual consumer innovators only. We have done this to highlight the contestable nature of their demand, and to emphasize that contestability can occur in markets for consumer goods. However, follow-on research could develop a similar model focused on or including user firms creating, for example, process innovations for their own use rather than for sale.

Second, we note that there are fields and markets in which some types of innovations originate only from innovating users—a situation with $s = 1$ in terms of our model. This is often the case, for example, with respect to the development of specialized techniques. Producers often find it impossible to profitably develop and market unprotectable techniques, and tend to leave that vital arena entirely or almost entirely to users (Hienerth 2016). In this paper we explored the importance of user innovation in markets that include producer innovation as well. However, further work could explore the nature of markets characterized by user innovation only.

Third, for simplicity, our model assumed that all innovating users will be able to benefit from a producer's investment in user innovation support, and that the producer will be able to observe the efforts of all innovating users and be able to reap any valuable spillovers. This is clearly not the case in practice—users will be differentially affected, and producers will not be able to observe or capture all spillovers generated by users. However, the same modeling logic and the same findings apply if our assumptions are true only for a subset of users.

Fourth, we assume that producers can choose the level of investment in support of innovating users that will maximize their profits. In the real world, users are independent actors who often have power to "push back" against producer plans and actions. They also can initiate user innovation activities in ways that producers do not expect. An example of investment in supporting user innovation not going according to producers' profit-maximizing plans is the case of Xara, a proprietary software company. In 2006, Xara invested in opening a large percentage of the source code of Xara Xtreme, a vector graphics package, as a way to invite user innovation. However, Xara did *not* open a small, commercially critical part of the source code. This omission caused a boycott among user programmers and, in the end, Xara yielded and opened more of the code than they would have preferred absent pressure from innovating users (Willis 2007).

It would be valuable and interesting for follow-on research to address situations such as the above. While in this paper we assumed that producers decide unilaterally to what extent they want to support and complement user innovation activities, we could think of a game in which innovating users can possess the power to determine the extent of user support, *s*, and potentially even the degree of complementarity, β. We expect that, in such a game, when the power to make both decisions lies with innovating users, they will pick higher levels of user support and complementarity than producers would. Unless users pick very high levels of *s*, this should lower producer profit but increase welfare overall. Future research could further explore this and also consider situations in which the decision power with regard to *s* and β is distributed between innovating users and producers.

Fifth, it is noteworthy that user innovators in our model receive no remuneration from producer firms. In the real world, successful user innovators sometimes receive payments for valuable contributions (such is the case with Lego and many app stores). Still, as a nationally representative survey in Finland shows, innovating users typically do freely reveal their innovations; our assumption of no payment is based on that situation (de Jong et al. 2015). In a different model, our variable *x* could be seen as the cost of user royalties to the firm, and implications for market outcomes could be explored.

Sixth, we have modeled producer support of user innovation as *increasing* the amount of time (or resources more generally) that users wish to spend on activities that benefit producers. Gamification of contributions and the setting-up of a user community were examples in point. It is also conceivable, however, that producer support, e.g., in the form of better tools, will enable users to *save time* while innovating. Such kinds of producer support could attract additional users to contribute, i.e., those who were previously non-innovators. This would endogenize the share of user innovators, σ, in a market, which we have taken to be exogenous in our model. It would be interesting for future research to explore the outcomes of this extended model, especially with regard to the optimal choice of producer strategy β.

Finally, our model treated all producers symmetrically, having all of them choose either a user-substituting or a user-complementing innovation strategy. Future research can usefully generalize from this limiting assumption. In the real world, we observe the coexistence of producers of both types. A key reason, we think, is that reorganizing and re-structuring R&D to exploit user-created innovation spillovers can be quite costly. Established firms with a legacy of producer-centric innovation will therefore be hesitant to switch, while new entrants without a commitment to the traditional model will likely find it economically more viable to choose the user-augmented innovation mode. Such constraints and switching costs could usefully be analyzed regarding their effects on strategic heterogeneity and firm and market-level outcomes. For instance, in markets with a growing share of user innovators, we should observe that new entrants and incumbents that are more flexible in organizing their R&D are more profitable.

Appendix A: Proof of the Choice-of-Mode Theorem

As noted, $z^U \leq z^P \rightarrow \Pi^P \geq \Pi^U$ and vice versa. Moreover, σ affects Π only through z and therefore we can study the impact of σ on Π by studying the impact of σ on z. We first show that at $\sigma = 0$, $z^U \leq z^P$, which establishes that at $\sigma = 0$ firms choose the P-mode. We then show that and if $0 < \beta < ½$, $z^U_\sigma \geq z^P_\sigma$, $\forall\ \sigma < \sigma_0$, with $\sigma_0 < 1$. For the first point, compare $\eta^U \tilde{b}^U$ and $\eta^P \tilde{b}^P = \xi$. At $\sigma = 0$, $\eta^U \leq \eta^P$ and $\tilde{b}^U = \tilde{b}^P = \xi$. As a result, $\sigma = 0$ implies $z^U \leq z^P$. For the second point, it is not difficult to see that $z^P_\sigma = 0$, and $z^U_\sigma = \tilde{b}^U(\eta^U_\sigma + \eta^U \tilde{b}^U_\sigma / \tilde{b}^U)$. Recall that $\eta^U_\sigma \leq 0$, and it is easy to see that $\tilde{b}^U_\sigma / \tilde{b}^U = \tau^\theta / [\sigma(\xi^\theta + \tau^\theta)]$. This expression is positive, and if $\theta < 1$, or $0 < \beta < ½$, it is very high when $\sigma \rightarrow 0$. Moreover, it declines as σ increases, and, given $\lambda \geq \mu$, η^U declines. All this implies that when σ is close to zero, $z^U_\sigma > 0$ because the positive value of $\eta^U \tilde{b}^U_\sigma / \tilde{b}^U$ outweighs the negative η^U_σ. As σ increases, z^U_σ declines, that is, $z^U_{\sigma\sigma} < 0$. This is because $z^U_{\sigma\sigma} = \tilde{b}^U_\sigma(\eta^U_\sigma + \eta^U \tilde{b}^U_\sigma / \tilde{b}^U) + \tilde{b}^U \partial(\eta^U + \eta^U \tilde{b}^U_\sigma / \tilde{b}^U) / \partial\sigma$, where based on what we have just said, the last derivative is negative. Then, when $\eta^U_\sigma + \eta^U \tilde{b}^U_\sigma / \tilde{b}^U = 0$, that is $z^U_\sigma = 0$, $z^U_{\sigma\sigma} < 0$, which in turn means that z^U reaches a maximum when $z^U_\sigma = 0$, and it then starts declining. This explains the shape of our curves in [the figure]. We have discussed in the text, and it is easy to see that when $\sigma = \lambda = \mu = 0$, then $\Pi^U = \Pi^P$. Then, as Π^U increases faster than Π^P as σ increases, a higher σ, with $\lambda = \mu = 0$, implies $\Pi^U > \Pi^P$. This establishes that a switch can take place if λ or μ are sufficiently small. Finally, differentiate $z^U - z^P$ with respect to λ or μ at $\sigma = \sigma^*$. We know that $z^U_\sigma - z^P_\sigma > 0$, and therefore the sign of σ^*_λ or σ^*_μ is the opposite of the sign of $z^U_\lambda - z^P_\lambda$ or $z^U_\mu - z^P_\mu$, which are both negative because λ and μ affect these expressions only through η^U. As a result σ* increases with λ or μ, and if λ or μ are too high the switch does not take place. QED

Appendix B: Proof of the Welfare Theorem

The strategy to prove this theorem is to show, first, that at σ^*, $W^U - W^P \geq 0$, and then that $W^U_\sigma - W^P_\sigma \geq 0$. Under the conditions of the choice-of-mode theorem, $\Pi^U_\sigma \geq \Pi^P_\sigma$. This means that at σ*, when the firms switch from the P- to the U-mode, welfare is higher in the U-mode, and for larger *σ*, welfare does not switch back to the P-mode.

To show that $W^U - W^P \geq 0$, evaluate $W^U - W^P = N(\Pi^U - \Pi^P) + (1 - \sigma)[\tfrac{1}{2}(1 - p^U + (1 - \mu)b^U)^2 + \mu b^U + \mu' h^U - \tfrac{1}{2}(1 - p^P + b^P)^2 - \mu' h^P] + \sigma[\tfrac{1}{2}(1 - p^U + (1 - \lambda)b^U)^2 + \lambda b^U + h^U - \tfrac{1}{2}(1 - p^P + b^P)^2 - h^P]$ at σ* where $\Pi^U - \Pi^P = 0$ and $z^U y^U = z^P y^P$, which implies $p^U = p^P$ and $z^U = z^P$. The latter equality implies $y^U = y^P$ and therefore $\eta^U b^U = b^P$. We can rewrite $W^U - W^P$ at σ* using all this information, suppressing for simplicity the superscript U, and rearranging terms, $W^U - W^P = (1 - \sigma)\tfrac{1}{2}[(1 - p + (1 - \mu)b)(1 - p + b) + \mu b(p - (1 - \mu)b)] + \sigma\, \tfrac{1}{2}[(1 - p + (1 - \lambda)b)(1 - p + b) + \lambda b(p - (1 - \lambda)b)] + \tfrac{1}{2}(1 - \eta)b - \tfrac{1}{2}(1 - p + \eta b)^2 + [(1 - \sigma)\mu' + \sigma](h^U - h^P)$. The terms in the first two square brackets are the number of users who buy times their surplus plus the number of users who do not buy times their surplus. This also explains why $0 \leq p - (1 - \mu)b \leq 1$ and $0 \leq p - (1 - \lambda)b \leq 1$. Beyond these boundaries the surplus of the non-innovating user is $\tfrac{1}{2} + \mu' b$ or $(\mu + \mu')b$, and $\tfrac{1}{2} + h$ or $\lambda b + h$ for the innovating users. As a result, $\mu b(p - (1 - \mu)b)$, $\lambda b(p - (1 - \lambda)b) \geq 0$ and it is easy to see that $h^U - h^P \geq 0$. Sum the first terms in the first two square brackets, weighed respectively by (1 – σ) and σ, and subtract $\tfrac{1}{2}(1 - p + \eta b)^2$. This yields $\tfrac{1}{2}[(1 - p + b)(1 - p + \eta b) - (1 - p + \eta b)(1 - p + \eta b)] \geq 0$ because $\eta \leq 1$ and $1 - p + \eta b \geq 0$ because η is a weighted average between (1 – μ) and (1 – λ), and $(p - (1 - \mu)b) \geq 0$. Since all the other terms in the expression for $W^U - W^P$ are non-negative, this establishes that at σ*, $W^U \geq W^P$.

The next step is to show that $W^U_\sigma \geq W^P_\sigma$. The expression for W^P is (10) using the specific expressions for PS, CS^{nui} and CS^{ui} with $\lambda = \mu = 0$ and *p*, *b* and *h* computed for the P-mode, which means that $x = 0$ and η, $\tilde{b}$ and *y* are obtained from the problem of the firm under the P-mode. It is easy to see that in this case σ does not affect η, $\tilde{b}$ and *y* and therefore $W^P_\sigma = (1 - \mu')h^P$. For the U-mode, $W^U_\sigma = N\Pi^U_\sigma + (1 - \sigma)[(1 - p + (1 - \mu)b)$

$(-p_\sigma + (1-\mu)b_\sigma) + \mu b_\sigma + \mu' h_\sigma^U] + \sigma[(1-p+(1-\lambda)b)(-p_\sigma + (1-\lambda)b_\sigma) + \lambda b_\sigma + h_\sigma^U] + [(1-p+(1-\lambda)b)^2 - (1-p+(1-\mu)\ b)^2]/2 + (\lambda-\mu)b + (1-\mu')h^U$, where apart from Π_σ^U and h^U we suppressed all the superscripts U. If $0 < \beta < ½$ and σ is close to zero, $\Pi_\sigma^U \geq 0$. Moreover, $h^U - h^P \geq 0$. Thus, to establish the sign of $W_\sigma^U - W_\sigma^P$ we need to show that all the other terms of the expression for W_σ^U are non-negative. Start with the last term. Rewrite the difference of squares as the product of the sum and difference of the two terms, and collect $(\lambda-\mu)b$. We obtain $(\lambda-\mu)b(p-(1-(\lambda+\mu)/2)b) \geq 0$ because $\lambda \geq \mu$ and $(p-(1-(\lambda+\mu)/2)b) = ½(p-(1-\lambda)b + p - (1-\mu)b)$ and we already established that $(p-(1-\mu)b)$, $\lambda b(p-(1-\lambda)b) \geq 0$. Finally, $(1-p+(1-\mu)b)(-p_\sigma + (1-\mu)b_\sigma) + \mu b_\sigma = (1-p+(1-\mu)b)(-p_\sigma + b_\sigma) + \mu b_\sigma(p-(1-\mu)b)$. We know that $(1-p+(1-\mu)b) \geq 0$, $(p-(1-\mu)b) \geq 0$, and $-p_\sigma + b_\sigma = -(\eta_\sigma b + \eta b_\sigma)/(N+1) + b_\sigma = -\eta_\sigma b/(N+1) + b_\sigma[1-\eta/(N+1)] \geq 0$. This is because $\eta_\sigma \leq 0$, $b_\sigma = \tilde{b}_\sigma y + \tilde{b} y_\sigma \geq 0$, and $1 - \eta/(N+1) > 0$ because $\eta \leq 1$. We obtain a similar result for the analogous term in λ. This establishes that $W_\sigma^U \geq W_\sigma^P$. QED

Appendix C: Proof of the Policy Theorem

Like in the previous theorem, the strategy to prove this theorem hinges on the fact that, as shown in the previous theorem, at $\sigma^* W^U - W^P \geq 0$, and then we study how $\Pi^U - \Pi^P$ and $W^U - W^P$ vary as we change γ or ξ. The logic is to check whether, under the conditions of the choice-of-mode theorem, changes in $\Pi^U - \Pi^P$ and $W^U - W^P$ go in the same direction.

In $W^U - W^P$ changes in γ do not affect any of the variables in the P-mode. They raise Π^U and h^U. The expression for W_γ^U is equivalent to W_σ^U in the proof of the previous theorem, without the last two terms and with subscripts γ instead of σ. Thus, to show that $W_\gamma^U - W_\gamma^P \geq 0$, we need to show that the second and third terms are non-negative. Analogously to the proof of the previous theorem, the second term can be written as $(1-p+(1-\mu)b)(-p_\gamma + b_\gamma) + \mu b_\gamma\ (p-(1-\mu)b)$. We know that $(1-p+(1-\mu)b) \geq 0$, $(p-(1-\mu)b) \geq 0$, and $-p_\gamma^U + b_\gamma^U = b_\gamma^U\left(1-\eta^U/(N+1)\right) \geq 0$. The same applies to the third term, which establishes that at σ^*, where $\Pi^U - \Pi^P = 0$ and $W^U - W^P \geq 0$, $W_\gamma^U + W_\gamma^P \geq 0$. This means that at σ^* increases in γ raise Π^U beyond Π^P, which induces firms to switch to the U-mode.

At the same time, welfare, which at σ* is higher in the U-mode, cannot turn to be smaller than in the P-mode.

Consider now increases in ξ. We first show that if β > β* with β* < ½, $\Pi_\xi^U - \Pi_\xi^P \leq 0$. To see this, $z_\xi^U - z_\xi^P = \eta^U\Psi - 1$ where $\Psi \equiv [1 + (\tau/\xi)^\theta]^{(1-\theta)/\theta}$. If 0 < θ < 1, or 0 < β < ½, $\eta^U \leq 1$ but Ψ ≥ 1. However, θ → 0 implies that Ψ becomes very large, and θ = 1 implies Ψ = 1. Moreover, Ψ declines monotonically with θ. Consider $\partial\log\Psi/\partial\theta = -\theta^{-2}\log\Psi' + [(1-\theta)/\theta][(\tau/\xi)^\theta/\Psi']\log(\tau/\xi)$, where $\Psi' \equiv [1 + (\tau/\xi)^\theta]$. This expression is negative because we study cases in which τ/ξ < 1. As a result, there is a threshold θ* < 1, or β* < ½, such that $\beta > \beta^* \rightarrow z_\xi^U - z_\xi^P < 0$, and vice versa. Thus, at σ = σ*, increases in ξ induce a switch to the P-mode.

To check for $W_\xi^U + W_\xi^P$, using the logic of the proof of the previous theorem we can write $W_\xi^U = N\Pi_\xi + (1-\sigma)[1 - p + (1-\mu)b)(-p_\xi + b_\xi) + \mu b_\xi(p + (1-\mu)b) + \mu' h_\xi] + \sigma[(1 - p + (1-\lambda)b)(-p_\xi + b_\xi + \lambda b_\xi(p + (1-\lambda)b) + h_\xi]$ where for simplicity we suppressed the superscripts U. The expression for W_ξ^P is the same with $\lambda = \mu = h_\xi = 0$ and the variables are all evaluated at the P-mode. Recall that, as noted in the proof of the previous theorem, at σ = σ*, $p^U = p^P$ and $\eta^U b^U = b^P$. Then, in $W_\xi^U - W_\xi^P$, the difference between $(1 - p + (1-\mu)b)(-p_\xi + b_\xi) + \mu b_\xi(p - (1-\mu)b)$ and the equivalent term in W_ξ^P, weighed by (1 – σ), and the difference between $(1 - p + (1-\lambda)b)(-p_\xi + b_\xi + \lambda b_\xi(p - (1-\lambda)b)$ and the equivalent term in W_ξ^P, weighed by σ, yields, after some algebra, $[(1-\eta) + \sigma(1-\sigma)(\lambda-\mu)^2 b]b_\xi \geq 0$. In addition, while h^P does not change with ξ, h^U increases in ξ. We conclude that at σ = σ*, where $W^U \geq W^P$, the sign of $W_\xi^U - W_\xi^P$ is ambiguous and can very well be positive. Since β > β* implies $\Pi_\xi^U - \Pi_\xi^P \leq 0$, it may be that a higher ξ induces firms to switch to the P-mode while welfare is still higher under the U-mode. QED

Notes

1. The ratio of the marginal products of $b = (T^\beta + Y^\beta)^{1/\beta}$ with respect to T and Y, is equal to $(Y/T)^{1-\beta}$. With σ small, T is small and therefore Y/T is likely to be larger than 1. As a result, a lower β makes the impact of T on b higher relative to the impact of Y on b. Since a higher σ makes T higher and Y lower because the optimal s increases, the condition 0 < β < ½ says that the contribution of the higher T on b has to be strong enough to compensate by a sizable amount the negative effect on b due to a lower Y.

2. The intuition is that increases in ξ have a direct positive impact on Y both in the P- and U-mode. In addition, in the U-mode increases in ξ reduce s, which raises Y and reduces T. However, a higher β generates a more pronounced drop in s relative to $(1 - s)$ because when complementarity is strong, the increase in Y does not produce a strong decline in T due to the feedback produced by complementarity. As a result, when complementarity is weak, increases in ξ produce a stronger increase in b in the P- than in the U-mode.

We hope you enjoyed reading this Open Access edition of *Free Innovation*.

If you would like to purchase a print copy, click here.

This is a section of doi:10.7551/mitpress/9382.001.0001

Free Innovation

By: Eric von Hippel

Citation:

Free Innovation
By: Eric von Hippel
DOI: 10.7551/mitpress/9382.001.0001
ISBN (electronic): 9780262335461
Publisher: The MIT Press
Published: 2024

The MIT Press

A hardcover edition of this text is available from The MIT Press. Copies may be ordered online at http://mitpress.mit.edu.

Set in Stone Sans and Stone Serif by Toppan Best-set Premedia Limited. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Names: von Hippel, Eric.
Title: Free innovation / Eric von Hippel.
Description: Cambridge, MA : MIT Press, 2016. | Includes bibliographical references and index.
Identifiers: LCCN 2016009390 | ISBN 9780262035217 (hardcover : alk. paper)
Subjects: LCSH: Technological innovations--Economic aspects. | Inventors. | Innovations.
Classification: LCC HC79.T4 H557 1988 | DDC 338/.064--dc23 LC record available at https://lccn.loc.gov/2016009390

10 9 8 7 6 5 4 3 2 1