Support Forums and Software Vendor's Pricing Strategy*

(Authors' names blinded for peer review)

January 22, 2020

Abstract

This research develops a micro-economic model to study the impact of support forums on the software industry. Motivated by the rapid growth of forums such as *Stack Exchange*, *Quora*, and *Yahoo! Answers*, we examine what impact they could have on a software vendor. Since these forums provide an alternative, free source for obtaining software support, one might suspect that this "shadow" competition eats into the profit of the vendor. We find that, while such a conclusion could be right at the initial stages of the forum's entry into the market, it may not be true in the later stages, especially if the forum generates enough momentum. This is because forums, besides playing a substitutive role, also have a complementary effect. At the initial stages, the substitutive role of the forum dominates. However, eventually, *if* and *when* the forum grows sufficiently, the complementary role takes control. In the mean time, the vendor can try to mitigate the immediate losses and find temporary relief by adopting a strategy of bundling support services with the product.

Keywords: Forums, stack exchange, software, support, bundling, pricing

1 Introduction

In recent times, online Q&A (questions and answers) forums, also known as support forums, have grown quite rapidly. These forums provide free online platforms where a user can post a wide variety of questions as well as answer others' questions. A support forum can be of two types. The *generic* ones, such as *Stack Exchange* and *Quora*, entertain questions on a wide variety of topics. In contrast, a *specific* forum, such as `office-forums.com` or `photoshopgurus.com/forum/`, deal with a specific product or topic; questions deemed irrelevant to the mission of the forum are deleted upon review.

Not surprisingly, product support communities are amongst the most popular ones within Q&A forums. In fact, for open source software products, Q&A forums are often the only vehicle for

*Draft only; do not quote or cite without permission.

product support. Even in the case of a commercial product, a consumer facing an issue with the product may pose a question in one of these product support forums; if another visitor to the forum knows the answer and decides to share it, the consumer may indeed end up getting a resolution to his issue. Since the entire conversation thread is stored in a database that is reorganized periodically for better efficiency, many other consumers with a similar issue may also gain from the content of this thread. As the forum gains momentum, consumers would start getting a larger share of their questions answered, and answered more satisfactorily. It is only natural that the entry and growth of this alternative source of free support would put a downward pressure on the consumer's willingness to pay (WTP) for vendor-provided support.

When a vendor faces such erosion in consumers' valuation for its services, several questions emerge, without immediate or known answers to any of them. How should the vendor react to the entry and growth of a support forum? Facing a "shadow" competition,[1] should the vendor resist, tolerate, or encourage such an entry? Can the vendor hedge against the loss that a forum might cause and, if so, how? Could there ever be a situation where the vendor actually benefits from the presence of such a forum? And, finally, what are the welfare implications of the entry and growth of a forum?

Answers to the questions above have important implications for the vendor's overall marketing strategy as well as broader connotations for public policy. Therefore, to answer these questions, we set up a parsimonious economic model comprising a software vendor with one software product and a service plan for providing after-sales support,[2] a mass of consumers who may purchase the software, with or without subscribing to the service plan, and a forum that provides free support, albeit of lower quality. We take into consideration that some consumers may not be savvy enough to use the forum effectively and may need vendor-provided support regardless. A key feature of our model is that we allow the vendor to strategically choose between: (i) selling the software and its associated support separately, and (ii) selling them together as a bundle.

A quick clarification is in order now. When we talk about product support being sold separately, we are *not* referring to the *nominal* services that always come for free with the purchase of a software

[1] We call it shadow competition because, in real competition, the competitor is also strategic. However, a forum is essentially free and does not set a price strategically.

[2] Although our story revolves around software, theoretically, our results apply to any information good (with zero or negligible marginal cost). In fact, our results extend to even physical goods, as long as the marginal cost is constant. Of course, in both cases, the concept of after-sales product support must be relevant for our results to make sense.

product. For example, almost all software vendors provide free security patches and (even) some feature updates, user account creation and management, and limited phone support for problems related to installation and user account management; such services incur negligible marginal costs and are always bundled with the product. The provisioning of any help/support beyond those basic, "free" services would be considered *premium*. Since it would have to employ dedicated/specialized call centers, the vendor is likely to face a non-zero marginal cost for these premium services. By support in this paper, we are alluding to this type of premium support plan that most software vendors—including Microsoft, Adobe, Oracle, Wolfram, and Intuit—sell.

Our modeling exercise provides some curious insights. It tells us that the entry and growth of a forum has a non-monotonic impact on the vendor's profit. Initially, the vendor's profit may decline as the forum grows in strength, but once the forum's strength reaches a threshold, the profit can indeed increase. In fact, in the extreme case in which the forum becomes a perfect substitute for vendor-provided support, the vendor's profit reaches its maximum, and the forum makes the vendor better off even when compared to the situation before the forum's entry. In essence, despite some early pains, the entry and growth of a support forum could eventually be quite beneficial to the vendor. We also find that the vendor can effectively hedge against the early losses by switching to a bundling strategy, even though such a strategy is never optimal absent the shadow competition from the forum.

Our analysis reveals that the driving force behind these surprising findings is a curious double role played by the forum. As discussed above, the forum poses as an imperfect substitute. At the same time, however, the forum also plays a complementary role by enhancing a consumer's WTP for the product itself. Although the forum's substitutive role dominates initially, its complementary role may eventually take control, especially if the forum gathers sufficient momentum. When that happens, the forum not only boosts private profits but also improves public welfare. In the mean time, the vendor can adopt a bundling strategy to put a floor on its immediate losses.

2 Practical Context: the Market for Support Services

Before we can analyze the impact of a forum's entry on the vendor's pricing strategy or interpret the meaning of our results, it is important to understand the practical context in which such an

entry is made.

A market for support services is common for many different types of products, digital or physical. Software is no exception. In order to help users when they are unable to use the software product for some reason, vendors offer professional support services, often at an additional fee. When required, a user may call a toll-free number and talk to an agent about the problem he faces to get a remedy. Consider Quicken, a popular financial management software for home use. It offers a *premier* plan that provides instant access to an agent over the phone who can help with a variety of problems, including those involving product installation, updates and patches, user registration, file management, and questions about a myriad of product features. Interestingly, Quicken also sells another version of the same product at a lower price-point under the *deluxe* plan, which provides minimal support to users and lacks the scope and timeliness of the premier plan. The deluxe plan is akin to selling only the software and is geared towards those consumers who may be tech-savvy and hence not interested in premier support.

This business strategy of offering support for a price is not limited to Quicken alone—it is a very common business model adopted by many software manufacturers. For example, Microsoft obtains about 38% of its revenue from services which include software-related support that it provides to the users of its products. Likewise, Intuit obtains 72% of its revenue from services alone. It is this type of software and support market that we are interested in. In particular, we consider the market for end-user commercial software products, such as Office, Mathematica, Photoshop, and Quicken, among several others. There are two main characteristics that define this market. First, it is usually monopolistic, with the vendor being the sole provider of support services—third-party support services, which are quite common for enterprise software products such as SAP or Peoplesoft, are rare for end-user software. Second, although the vendor incurs no marginal cost for the product itself, the marginal cost for support services is quite significant (Jabr et al. 2014). Examples of this type of software abound, a few of which are listed in Table 1.

It is apparent from Table 1 that the market for end-user software has several major players. Among them, Microsoft alone sold personal computer software worth $45.7 billion in 2019, a large part of which is from its Office suite. Similarly, Adobe's 2019 revenue was $10.6 billion, all of which came from end-user products. Clearly, this market is too large and too significant not to demand a better scrutiny. Interestingly, as far as bundling is concerned, there is no clear consensus among

Table 1: Examples of Software and Support Markets

Product	Vendor	Forums
Product and Support Sold Separately		
Photoshop	Adobe	community.adobe.com* graphicdesign.stackexchange.com photoshopgurus.com/forum/ tek-tips.com/threadminder.cfm?pid=229
Acrobat	Adobe	forums.adobe.com* tek-tips.com/threadminder.cfm?pid=223
Mathematica	Wolfram	community.wolfram.com* mathematica.stackexchange.com
Quicken	Quicken	community.quicken.com*
Office	Microsoft	office-forums.com* tek-tips.com/threadminder.cfm?pid=68 msofficeforums.com*
PowerDirector	Cyberlink	forum.cyberlink.com* goproforums.com/threads/cyberlink-powerdirector-11.1578/
Product and Support Sold as a Bundle		
Matlab	Mathwork	mathworks.com/matlabcentral* forums.codeguru.com/showthread.php?396124-Matlab-Help-Forum stackoverflow.com/questions/tagged/matlab youth4work.com/Talent/Matlab/Forum
Scientific Word	MacKichan	forum.mackichan.com*
CorelCAD/ CorelDRAW	Corel	community.coreldraw.com/talk/* forum.corel.com/EN/* tek-tips.com/threadminder.cfm?pid=238 stackoverflow.com/questions/tagged/coreldraw
TurboTax	Intuit	quickbooks.intuit.com/learn-support/*
Camtasia	Techsmith	feedback.techsmith.com/techsmith* camtasia.wonderhowto.com/forum/ steamcommunity.com/app/253050/discussions/

*Forums marked with an asterisk represent forums that are maintained by the vendor itself.

the vendors—some vendors bundle support with the product, while others sell them separately. In fact, there are instances where the same vendor bundles support for one product while selling separately for another.

Table 1 also sheds significant light on the status of online forums related to software products. First, support forums are quite common for end-user software products. Every product listed in the table has at least one online support forum. Second, even in cases where support is bundled with the product, there could be third-party forums. Consider, for example, the cases of Matlab and CorelCAD. Since support for each of these products is already bundled, its users do not have to pay anything extra when they seek support, making a related forum much less attractive to a user. Yet, there exist several forums for these products, most of which are independent of the vendor. Finally, in certain instances (such as Quicken and Scientific Word), there is only one forum, and it

is maintained by the vendor; in these cases, an independent forum does not exist.

Interestingly, our modeling experiment can explain many of the practical observations about these markets. For example, the ubiquitous presence of vendor-supported forums can be explained from our result that, eventually, a forum can enhance a vendor's profit. It is for the same reason that we see vendors participating in independent forums as well. Similarly, that some vendors bundle support with the product while others do not is consistent with our findings. According to our results, vendors should bundle only when the forum's growth is moderate and should sell the product and support separately at the two extremes. Since support forums for different software products may be at different stages of growth, they are likely to have different levels of penetration and effectiveness. The vendors' bundling strategies should naturally be different.

3 Literature Review

Our paper contributes to the growing literature on online communities. This literature has become so vast in recent years that it is difficult to provide a comprehensive review here. Among the recent works of interest are the ones that involve examination of online communities engaged in software support (e.g., Huang and Zhang 2016, Huang et al. 2018, Jabr et al. 2014, Lu et al. 2017) as well as those that examine question-and-answer (Q&A) forums (e.g., Khansa et al. 2015, Lappas et al. 2017, Zhao et al. 2016). Most of these papers investigate the incentives of the community members to contribute as well as the optimal design from the platform owner's perspective.

Briefly, the visitors to a forum can be thought of as a continuous spectrum. At one end of this spectrum, there are *questioners*, users who only ask questions but do not contribute with solutions (Jabr et al. 2014). At the other, there are *solvers* or experts, those who answer questions but do not ask any. According to Huang et al. (2018, p.214), these experts are "typically early adopters of products or services, have up-to-date knowledge with regard to the practice in question, and are the most active knowledge contributors to virtual communities." Interestingly, solvers may as well be employees of the software vendor, serving as seeds of knowledge at its behest.

Prior research has also looked at the possible reasons an expert contributes to a forum. There are several, recognition among peers as a problem solver being an important one (Jabr et al. 2014). Such recognition is not just satisfying, but it can also translate to material benefits, such

as a promotion within the employer organization for helping consumers and strengthening the employer's brand, or job offers from firms in need for similar expertise (Huang and Zhang 2016). Of course, there may also be some experts who contribute for reasons that are purely altruistic.

Now, experts need not be the only ones who contribute to a forum. This is because the forum operator puts various incentives in place to draw contribution from as many users as possible. Such incentives can come in the form of badges awarded to top contributors, based on the number of answers or the number of users who found the answers helpful (Jabr et al. 2014). The experts might serve as the seed in the initial stages of the forum, but motivated by these incentives and/or altruistic reasons, a user, who begins as a questioner, may eventually start passing on the knowledge he gains (from the experts) to others in the community. Thus, a questioner can gradually become a solver. More broadly, "the extent of users' knowledge contributions to the community depends on their domain expertise and ability to discover existing knowledge, conduct independent investigations and experiments, recombine ideas shared by peer members, and find novel solutions" (Huang et al. 2018, p.217). The net result of this complex knowledge co-creation and diffusion is that, between the extremes of questioners and solvers, there are likely many who both ask and answer questions. In fact, it is this broad engagement from users with different levels of aptitude and knowledge that makes the forum a vibrant one and marks its eventual success.

Another related point is that, as a readily accessible knowledge base, forums can sometimes provide the necessary help in a manner that is more efficient and timely than even professional support. This is because a forum typically organizes and stores all the past questions along with their posted answers which a user can browse or search on demand. Naturally, if the issue is a common one, the user would very likely find a solution even without having to ask the question explicitly and wait for someone to answer it. In fact, this way, a forum may have users who are not even registered, particularly if the forum does not require users to log on to browse the available content. Such users would obviously fall in the category of pure questioners.[3] In our model, we do not delve into exactly how a user gets a problem resolved—we simply treat forums as a "black box" that can be used as an alternative source of support. This simplification is not a limitation of our work, though, because our objective is not to analyze how various knowledge

[3]Stack Overflow, one of the popular forums, has over ten million registerd users, but the number of unique visitors that it attracts every month exceeds forty million. Clearly, there are millions of unregistered questioners, who are attracted by the huge inventory of more than sixteen million questions, many already answered by multiple solvers.

communities grow or operate. Rather, our focus is on how their entry and growth might impact a vendor's pricing strategy. Put another way, our focus is on the vendor and not on the forum per se. We complement earlier works by identifying how a vendor can adjust its pricing strategy to best leverage the presence of such a forum. We also explain how the vendor's strategy ought to evolve when user participation and contributions to the forum increase.

Modeling-wise, our research is related to the literature on asymmetric bundling, which considers the problem of bundling an ancillary service or add-in, with a physical or software product (Dewan and Freimer 2003, Pang and Etzion 2012). The bundling problem is called asymmetric because the product has standalone value but the ancillary component does not. To elaborate, Dewan and Freimer (2003) consider the problem of bundling a spell-checker with a word-processor software. In their context, the spell-checker is of little use to a consumer who has not purchased the word-processor, but the converse is clearly not true. Thus, their setting is different from the traditional bundling literature where each product has its own standalone value (Adam and Yellen 1976, Bakos and Brynjolfsson 2000). In our context, the support offered by the vendor is essentially an ancillary component, which is valuable only to consumers who have purchased the software.

The notion of asymmetric bundling is also related to that of versioning (Bhargava and Choudhary 2008). After all, in our context, a consumer has only two options, buying only the software or buying both the product and support. This is essentially a vertically differentiated setup, where buying both is equivalent to acquiring the higher quality version, whereas buying only the product, the same as getting the lower quality one. Despite our context being one of asymmetric bundling—and hence related to versioning—how does our work contribute to the existing literature in these streams? First, in our context, the ancillary component—support—has an imperfect but a free substitute in the form of a support forum. Second, and more importantly, the support forum not only acts as a free substitute for vendor-provided support, but it can also enhance a consumer's valuation for the product itself, thereby acting as a complement as well. These twin roles played by the forum—a substitute and a complement at the same time—are new and has not been examined in prior research.

To elaborate further, in the presence of a forum, savvy consumers who are comfortable with getting their support from the forum will value the software more than the non-savvy ones who depend critically on vendor-provided support. At the same time, however, the savvy users will

value vendor-provided support less vis-à-vis the non-savvy. This negative correlation creates an opportunity for the vendor to use bundling, an opportunity not present in the no-forum scenario. Viewed another away, the vendor's problem is one of straightforward versioning if there is no forum, but bundling creeps in as a possibility when there is one. The key question then is: How does the efficacy of bundling depend on the adoption and effectiveness of the forum? This is an issue that we discuss in detail, offering new insights and enriching prior literature in the process.

4 The Basic Setup

We consider two types of players—a monopoly software vendor and a group of consumers whose mass is normalized to unity. Consumers are heterogeneous in their valuation for the product offered by the vendor:

Assumption 1 *Consumers are indexed by v, which represents their WTP for the product alone; v is distributed uniformly over $[0, 1]$. A consumer knows his v, but the vendor knows only the distribution.*

We expect a consumer's overall WTP in the presence of vendor-provided support to be higher; this increase in WTP, denoted $w(v)$, ought to be an increasing function of v satisfying $w(0) = 0$ and $w(v) < v$; we assume a linear form for the sake of exposition:

Assumption 2 *Vendor-provided support enhances consumer v's overall WTP by $w(v) = \beta v$, where $\beta \in (0, 1)$ is a constant.*

We assume the following cost structure for the vendor:

Assumption 3 *The vendor incurs a zero marginal cost for the product and a cost of $c > 0$ per user for support.*

Now, the vendor can either sell the product and support as a bundle. or it can sell them separately. If the product and support are sold separately by the vendor, at prices $p > 0$ and $s > 0$ respectively, consumer v would derive a net benefit of $(v - p)$ from the product alone. So, initially, the consumer will buy the product if $v \geq p$. At this time, however, the consumer does not have sufficient information to decide whether he would also need the vendor-provided support. This is

because the *reliability* of the product is private information to the vendor,[4] so whether any support is needed or not is initially unknown to the consumer. Furthermore, the need is likely to depend on the specific configuration and features being used by the consumer. In summary, only after using the product ample times can the consumer truly assess if he has a need for vendor-provided support.

Let r be the probability that such a need would arise; r can also be viewed as the fraction of consumers who would feel the need for professional support after using the product. Although unknown to consumers, r can be estimated by the vendor from the results of its internal testing of the product and its reliability, ρ. In particular, r ought to be a decreasing function of ρ, that is, the vendor can lower the consumers' need for support by making *additional* investments towards ρ; we assume a simple linear form:

Assumption 4 *The probability of a consumer needing support after use is $r = 1 - \rho$. In order to obtain a reliability level of $\rho \in (0, 1)$, the vendor has to invest $\kappa(\rho)$, where $\kappa(0) = 0$ and $\kappa'(\rho) \geq 0$.*

To keep the model parsimonious, we assume that the consumer first uses the product and derives the base value of v. It is only after using it for a period that he realizes his actual need for support. If it turns out that he does not need any additional support after all, he can continue using the product without support and thereby enjoy the additional value of βv. However, if a need indeed arises, he now has to weigh the option of paying for support with that of forgoing it. If he forgoes, he only gets a net utility of $v - p > 0$ from the product. However, on top of this, he obtains an additional amount of $\beta v - s$, if he purchases support too. Therefore, the consumer purchases support if and only if $\beta v \geq s$. Given this model of consumer behavior, the total profit to the vendor is simply $\pi = p(1 - p) + (1 - \rho)(s - c)\left(1 - \frac{s}{\beta}\right) - \kappa(\rho)$. Maximizing π over ρ, p, and s, subject to the constraint $p \leq \frac{s}{\beta}$, we get:

Proposition 1 *When selling the product and its support separately, the vendor sets $\rho^* = 0$, $p^* = \frac{1}{2}$, and $s^* = \frac{\beta(1+\chi)}{2}$, and makes a total profit of $\pi^* = \frac{1+\beta(1-\chi)^2}{4}$, where $\chi = \frac{c}{\beta}$ is the normalized marginal cost incurred by the vendor for providing support.*

[4] Is reliability of a product related to its quality? Broadly speaking, quality is a multi-dimensional concept and can mean many things within a context. In the literature, quality has often been viewed as the product's feature-richness that enhances a consumer's WTP for it. To the extent that the product's appeal to a consumer depends on its reliability, reliability can indeed be viewed as a part of quality. Consistent with prior research, here we are using the term reliability to refer to the product's robustness that determines consumers' future need for any help or support and not the product's overall appeal to them.

The boundary solution $\rho^* = 0$ must not be taken to mean that the overall reliability level of the product is actually zero. Of course, a vendor has to maintain a reasonable level of reliability, if only for the reputation of the product or brand. Within the context of our positive modeling experiment, this simply means that the vendor would maintain a nominal reliability level that is barely enough to sell its product at the price it wants to charge, but it would not make any *additional* investments towards reliability that actually drives its support market down. Therefore, Proposition 1 tells us that, in order to induce as many consumers as possible to purchase support, the vendor intentionally keeps the reliability level low and makes little additional investment in this regard, thereby ensuring that a substantial fraction of consumers feel the need for professional support after using the product, that is, r^* is sufficiently high.

Now, there is another way the vendor can compel every consumer to buy support—it can simply bundle the support plan with the product and sell the bundle for a combined price of p_b. If indeed the vendor chooses to do so, the net benefit to consumer v would be $v(1+\beta) - p_b$, which must be greater than zero for the consumer to buy the bundle. The profit from the bundle would then be $\pi_b = (p_b - c)\left(1 - \frac{p_b}{1+\beta}\right)$, maximizing which we get:

Proposition 2 *When selling the product and its support as a bundle, the vendor sets the price at* $p_b^* = \frac{1+\beta(1+\chi)}{2}$ *and makes a total profit of* $\pi_b^* = \frac{(1+\beta(1-\chi))^2}{4(1+\beta)}$*, where* $\chi = \frac{c}{\beta}$ *as before.*

Comparing π_b^* in Proposition 2 with π^* in Proposition 1, we can see that $\pi_b^* - \pi^* = -\frac{\beta\chi^2}{4(1+\beta)} < 0$ for all $\beta, \chi \in (0,1)$, implying:

Corollary 1 *Irrespective of* β *and* χ*, it is not optimal to bundle the product and its support.*

The result in Corollary 1 is consistent with its counterpart in the versioning literature. When $\chi = 0$, bundling and selling separately yield identical profits, but bundling is strictly suboptimal for all $\chi > 0$. Similarly, at $\chi = 0$, the vendor is indifferent between selling one or two versions, but for all $\chi > 0$, selling just one version is suboptimal (Bhargava and Choudhary 2008).

5 Entry of a Support Forum

We now analyze how the vendor's pricing strategy would be impacted by the entry of a support forum. We represent the support forum by its *penetration*, λ, and *effectiveness*, μ; $\lambda, \mu \in (0,1)$.

The first parameter, λ, denotes the fraction of consumers who are savvy enough to be able to use the forum—when the forum's penetration or reach is high, so would be λ. The second parameter, μ, is the forum's perceived effectiveness; it denotes the quality of support obtained from the forum vis-à-vis that sold by the vendor. The quality difference could arise from several reasons. First, the forum may not offer fixes of the same quality, and in some cases may only offer inferior workarounds. Further, the support received may not be as prompt as that received from the vendor. Finally, the forum may not have answers to all the questions a user has. Formally, the presence of the forum creates another level of heterogeneity among consumers:

Assumption 5 *The presence of a forum separates consumers into two segments: (i) a λ fraction of forum-savvy (S) ones who are able to use the forum to get a μ fraction of their support-related needs met, and (ii) the remaining non-savvy (NS) fraction who cannot use the forum at all.*

Two observations are in order before we proceed further. First, consumers are unlikely to know the true effectiveness of the forum, but they should still be able to gauge its effectiveness in an expected sense; thus, μ represents that expected effectiveness and is the same for all consumers. Second, it is possible that μ depends on λ. After all, it is only reasonable that, as the forum attracts more people, its knowledge-base would grow as well, thereby enhancing its ability to resolve more issues. We do recognize this dependence and will discuss it in Section 7.1 in details. For now, we assume that these two dimensions are orthogonal.

If a consumer does not need support after using the product, he goes on to enjoy an additional value of βv. However, if he needs support and if he is savvy, he would get his issues resolved by the forum with a probability of μ. Therefore, he would derive an expected additional benefit of $\beta\mu v$, which is, of course, less than βv for all $\mu \in (0, 1)$. Thus, an immediate implication of Assumption 5 is that, irrespective of his need for support, the savvy user is assured of a minimum additional benefit of $\beta\mu v$ the moment he purchases the product, and his real valuation for the product should then increase to $v(1 + \beta\mu)$. Clearly, he would buy the product if $v(1 + \beta\mu) \geq p$, and if a need for the vendor-provided support arises later, he would buy it as well if $\beta v(1 - \mu) \geq s$. In contrast, if he is not forum-savvy, he would derive no benefit from the forum. Therefore, his behavior would remain exactly as described in Section 4.

Depending on the choice of p and s, theoretically, four cases are possible:

- **Case 1:** $p \leq \frac{s}{\beta}$ and $\frac{s}{\beta(1-\mu)} < 1$. In this case, irrespective of whether savvy or not, there are some consumers who buy just the product and do not consider buying support at all, while others buy support if a need arises.

- **Case 2:** $p \leq \frac{s}{\beta}$ and $\frac{s}{\beta(1-\mu)} \geq 1$. In this case, only some consumers from the non-savvy group buy support when the need arises, but none from the savvy group considers buying support.

- **Case 3:** $p > \frac{s}{\beta}$ and $\frac{p}{1+\beta\mu} < \frac{s}{\beta(1-\mu)} < 1$. In this case, all consumers from the non-savvy group buy support when the need arises, while only some in the savvy group do so.

- **Case 4:** $p > \frac{s}{\beta}$ and either $\frac{p}{1+\beta\mu} \geq \frac{s}{\beta(1-\mu)}$ or $\frac{s}{\beta(1-\mu)} \geq 1$. In this case, all consumers from the non-savvy group buy support when needed. The savvy consumers, in contrast, either all buy support when needed, $\frac{p}{1+\beta\mu} \geq \frac{s}{\beta(1-\mu)}$, or no one does, $\frac{s}{\beta(1-\mu)} \geq 1$.

Lemma 1 *Case 4 cannot occur in equilibrium.*

Lemma 1 tells us that only the first three market configurations are relevant in equilibrium; they are all depicted in Figure 1. Please note that each of these three configurations must always satisfy $\frac{p}{1+\beta\mu} < p$ and $\frac{s}{\beta} < \frac{s}{\beta(1-\mu)}$, because $\beta, \mu \in (0,1)$. It is clear from Figure 1 that the expression for profit from product sales alone is the same in each case: $\pi_p = \lambda p\left(1 - \frac{p}{1+\beta\mu}\right) + (1-\lambda)p(1-p)$; the profit from support, however, is different, leading to the following expressions for the vendor's total profit:

$$\pi = \begin{cases} \pi_1 = \pi_p + (1-\rho)(s-c)\left(\lambda\left(1 - \frac{s}{\beta(1-\mu)}\right) + (1-\lambda)\left(1 - \frac{s}{\beta}\right)\right) - \kappa(\rho), & \text{Case 1,} \\ \pi_2 = \pi_p + (1-\rho)(s-c)(1-\lambda)\left(1 - \frac{s}{\beta}\right) - \kappa(\rho), & \text{Case 2,} \\ \pi_3 = \pi_p + (1-\rho)(s-c)\left(\lambda\left(1 - \frac{s}{\beta(1-\mu)}\right) + (1-\lambda)(1-p)\right) - \kappa(\rho), & \text{Case 3.} \end{cases} \tag{1}$$

Maximizing this profit over ρ, p, and s, we find that, in equilibrium, the parameter space gets partitioned into five regions—one interior solution for each of the three relevant cases and two corner (limiting) solutions where the vendor holds p at its limiting value of $\frac{s}{\beta}$. Formally:

Proposition 3 *Let $\chi = \frac{c}{\beta}$ as before. Then, $\rho^* = 0$, and the equilibrium occurs in one of the following five regions within the parameter space:*

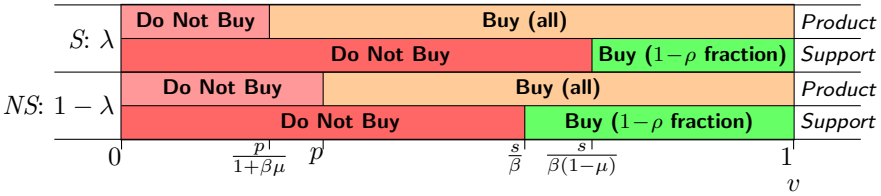

(a) **Case 1**: $p \leq \frac{s}{\beta}$ and $\frac{s}{\beta(1-\mu)} < 1$

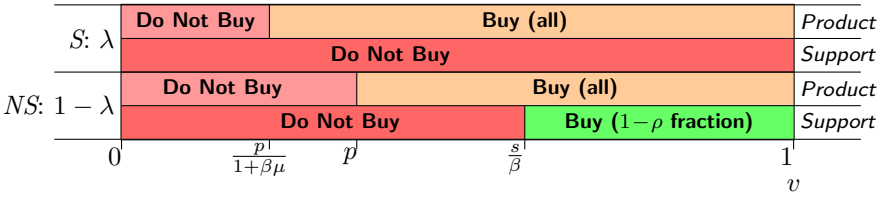

(b) **Case 2**: $p \leq \frac{s}{\beta}$ and $\frac{s}{\beta(1-\mu)} \geq 1$

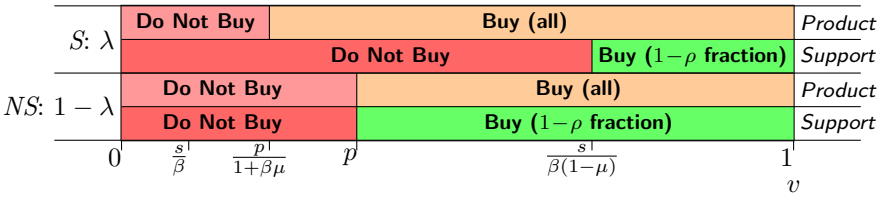

(c) **Case 3**: $p > \frac{s}{\beta}$ and $\frac{p}{1+\beta\mu} < \frac{s}{\beta(1-\mu)} < 1$

Figure 1: Possible Market Configurations

- **Region 1:** *In this region, the vendor charges* $p_1^* = \frac{1+\beta\mu}{2(1+\beta\mu(1-\lambda))}$ *and* $s_1^* = \frac{\beta}{2}\left(\frac{1-\mu}{1-\mu(1-\lambda)} + \chi\right)$, *such that* $p_1^* \leq \frac{s_1^*}{\beta}$ *and* $\frac{s_1^*}{\beta(1-\mu)} < 1$.

- **Region 1L:** *This region occurs as a corner solution of Case 1. In this region, the vendor charges* $p_{1\text{L}}^* = \frac{(1+\beta\mu)((1-\mu)(1+\beta(1+\chi))+\beta\chi\lambda\mu)}{2(1+\beta)(1-\mu(1-\beta(1-\mu(1-\lambda))))}$ *and sets* $s_{1\text{L}}^* = \beta p_{1\text{L}}^*$; *the solution is valid only if* $\frac{s_{1\text{L}}^*}{\beta(1-\mu)} < 1$.

- **Region 2:** *In this region, the vendor charges* $p_2^* = \frac{1+\beta\mu}{2(1+\beta\mu(1-\lambda))}$ *and* $s_2^* = \frac{\beta(1+\chi)}{2}$, *such that* $p_2^* \leq \frac{s_2^*}{\beta}$ *and* $\frac{s_2^*}{\beta(1-\mu)} \geq 1$.

- **Region 2L:** *This region occurs as a corner solution of Case 2. In this region, the vendor charges* $p_{2\text{L}}^* = \frac{(1+\beta\mu)(1+\beta(1+\chi)(1-\lambda))}{2(1+\beta(1-\lambda)(1+\mu(1+\beta)))}$ *and sets* $s_{2\text{L}}^* = \beta p_{2\text{L}}^*$. *A valid solution in this region must abide by* $\frac{s_{2\text{L}}^*}{\beta(1-\mu)} \geq 1$.

- **Region 3:** *In this region, the vendor charges* $p_3^* = \frac{(1+\beta\mu)(\lambda(2+\beta\chi(1-\lambda))-\beta(1-\lambda)(1-\mu))}{4\lambda-\beta^2\mu(1-\mu)(1-\lambda)^2-\beta(1-\lambda)(1-\lambda-\mu-3\lambda\mu)}$ *and*

$s_3^* = \frac{\beta(2\chi\lambda+\beta\mu(1-\lambda)(1-\mu)(1-\beta\chi(1-\lambda))+((1+\lambda)(1-\mu)-\beta\chi(1-\lambda)(1-\lambda-\mu-\lambda\mu)))}{4\lambda-\beta^2\mu(1-\mu)(1-\lambda)^2-\beta(1-\lambda)(1-\lambda-\mu-3\lambda\mu)}$, *such that* $p_3^* > \frac{s_3^*}{\beta}$ *and* $\frac{s_3^*}{\beta(1-\mu)} < 1$.

It is comforting to see that the results in Proposition 3 for Regions 1 and 2 reduce to Proposition 1 when either $\lambda \to 0$ or $\mu \to 0$. Proposition 3 is illustrated in Figure 2, where the entire (λ, μ)-space is partitioned into the five regions. There are several interesting observations that can

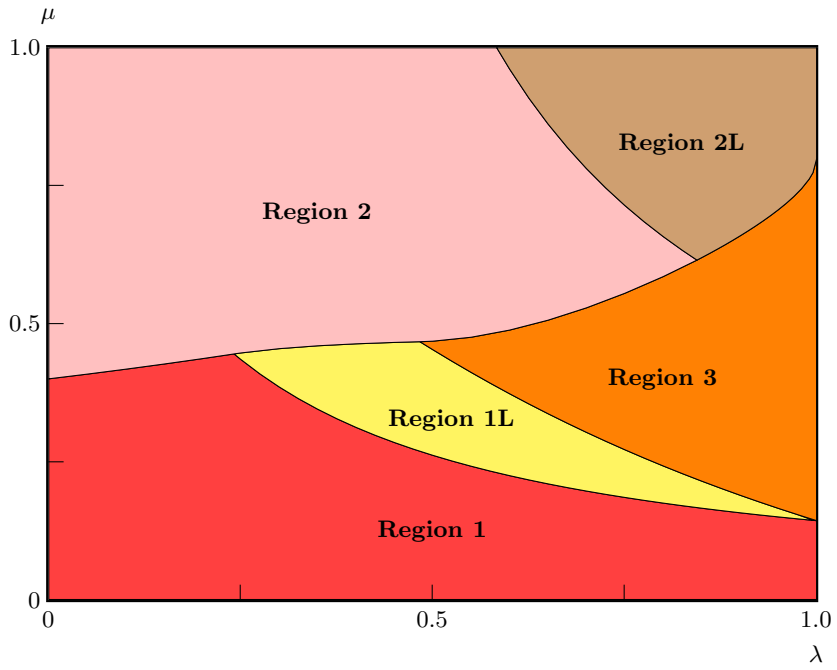

Figure 2: Different Equilibium Regions as Partitions of the (λ, μ) Space ($\beta = 0.4$, $\chi = 0.2$)

be made from Proposition 3 and Figure 2. First, in Regions 1 and 2, the price of the software product is independent of χ, the normalized support cost. These two regions, according to Figure 2, occur for low as well as high values of μ, provided λ is not too large. When either λ or μ is small, the forum poses little competitive threat to the vendor, and the vendor's pricing strategy is somewhat similar to the one it adopts in the absence of the forum—it takes the marginal support cost into consideration only when setting a price for the support, but the price of the product remains largely independent of the cost. As λ and μ increase—that is, as the forum gains in penetration and effectiveness—the forum starts substituting for vendor-provided support and exerts competitive pressure on the vendor. This is why in Regions 1L, 2L, and 3, which occur for moderate or large values of λ and μ, the two prices become a bit more tightly coupled, so much so

that the marginal cost of support now shows up in both the prices.

Interestingly, the forum plays two opposing roles at the same time. It acts as an (imperfect) substitute for the vendor-provided support, thereby putting competitive pressure on the vendor. However, it also acts as a complement to the underlying software product by elevating the WTP of the savvy consumers. When λ and μ are both moderate, or when either one is low, the forum's substitutive role dominates, but when they are both high, the complementary role takes over. Recognizing these twin roles played by the forum is quite necessary to grasp the full import of our subsequent results.

In order to get further insights about the equilibrium, we estimate the vendor's optimal profit in Region i, $i \in \{1, 1\text{L}, 2, 2\text{L}, 3\}$, by substituting the optimal prices in these regions from Proposition 3 into the respective profit expressions in (1). Figure 3 illustrates how this profit changes with the forum's penetration and effectiveness. A closer look at the profit plot in Figure 3 reveals a

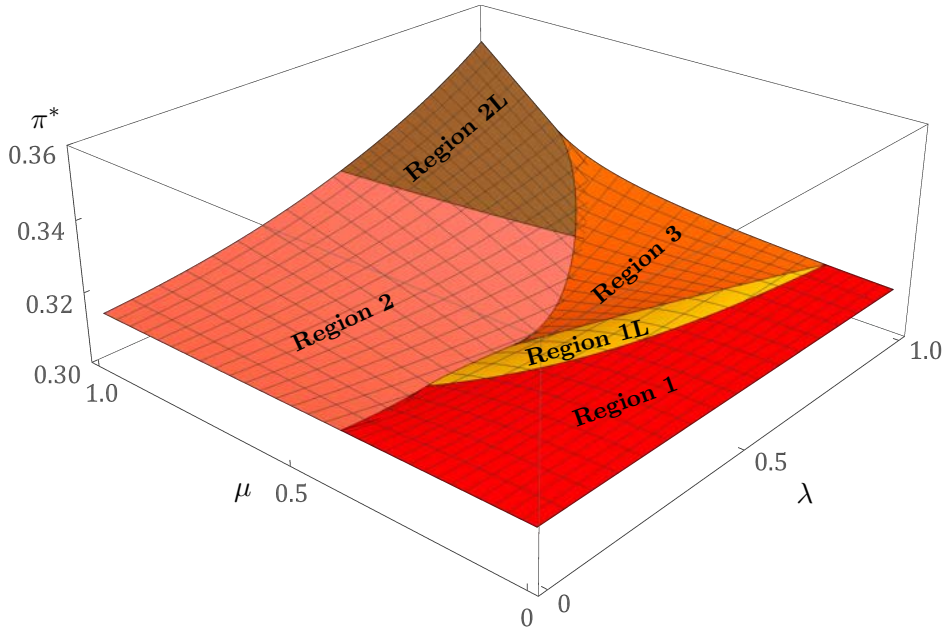

Figure 3: Vendor's Profit in Equilibrium ($\beta = 0.4$, $\chi = 0.2$)

curious observation—the vendor's profit is non-monotonic in both λ and μ. Initially, this might come across as somewhat counterintuitive—since the forum acts as a substitute and reduces some consumers' WTP for support, one might expect the vendor's profit to decrease when the forum gathers momentum, either in penetration or in effectiveness. However, one could also argue in the

opposite direction. Since the forum acts as a complement to the product and increases certain consumers' WTP for it, one could also expect the profit to increase with λ and μ. It turns out that the influence of the two opposing roles of the forum is not uniform over the entire parameter space. One dominates the other in some parts of the parameter space and, at the same time, gets dominated in the remaining parts, leading to a non-monotonic behavior.

More specifically, the non-monotonic behavior manifests itself in such a way that the vendor is better off when either λ or μ is high or low but, for moderate values of λ and μ, the vendor's profit experiences a significant dip. Essentially, the substitutive role of the forum takes hold for low and moderate values of λ and μ. Of course, when λ or μ is small, the forum lacks sufficient penetration or effectiveness, and its ability to act as a substitute is quite limited, which allows the vendor to still enjoy a handsome profit. However, as λ and μ start to increase, the forum gains more momentum, and the vendor's profit takes a downward turn in the face of this increased "shadow" competition. This leads to a visible dip in the profit in Figure 3 at moderate values of λ and μ. At even higher levels of λ and μ, however, the complementary role of the forum takes over, and the vendor's profit experiences an upsurge once again. This observation is quite general and applies to a large portion of the parameter space; as long as the support cost, χ, is not too high—that is, as long as the vendor enjoys a sufficient margin from selling support—the vendor's profit nosedives at moderate values of λ and μ.

The vendor's profit is actually maximized at $\lambda = \mu = 1$, that is, when the forum becomes a perfect substitute for vendor-provided support. This is because, just by pricing the product higher, the vendor is able to extract every bit of additional consumer surplus generated by the forum. This way, despite its substitutive role, a strong support forum may actually be to the vendor's advantage. This is perhaps why so many vendors actually monitor popular online forums and encourage their own support staff to contribute to these forums with tips, remedies, and work-arounds. It is, however, important to recognize that a popular and effective forum does not take shape overnight. These forums are often characterized by rich network effects (Parker and Van Alstyne 2005), and they may take some time before gathering the initial momentum. Therefore, while a strong forum can enhance a vendor's profit eventually, its immediate implications could be quite adverse, especially if the vendor is making good profit from the support provision already.

When the forum has just about enough strength to provide a shadow competition to vendor-provided support but is still sufficiently weak so that its complementary role is subdued, the vendor's profit may go down significantly.

6 Bundling as a Strategy

Given the adverse effects of the forum on the vendor's immediate profits, we now ask: Can the vendor do anything to avoid this predicament? More specifically, can the vendor use bundling as a strategy to hedge against this loss in profit? Of course, from Corollary 1, we know that, absent a support forum, it is suboptimal for the vendor to sell the product and its support as a bundle. However, now in the presence of a support forum, given the dwindling profit for the vendor at moderate values of λ and μ, could bundling actually become an optimal strategy for the vendor? Answering this question would require us to compare the profit in each region in Figure 3 with what the vendor can earn through bundling.

Interestingly, if the vendor were to choose bundling in the presence of the forum, its profit would still be given by $\pi_b^* = \frac{(1+\beta(1-\chi))^2}{4(1+\beta)}$; see Proposition 2. This is because, when the product and its support are bundled together, consumer v, regardless of his savvy, has a WTP of $v(1+\beta)$ for the bundle. A direct implication of this fact is that the bundling profit, π_b^*, is independent of λ and μ. When we compare this π_b^* with the profits in Figure 3, we find that an interior region emerges in the parameter space where, in fact, bundling can now be optimal. We illustrate this in Figure 4; the associated profit plot is presented in Figure 5.

To formally characterize the boundaries separating the bundling region, we perform pairwise profit comparison to obtain:

Lemma 2 *Let $\chi = \frac{c}{\beta}$ be as above. Further, for all $i \in \{1, 1\text{L}, 2, 2\text{L}, 3\}$, let $h_i(\lambda, \mu)$ be the solution of $\pi_i^* = \pi_b^*$. Then, the vendor's strategy can be summarized as:*

(i) Case 1 is valid if $\chi \geq g_1(\lambda, \mu) = \frac{1+\beta\mu}{1+\beta\mu(1-\lambda)} - \frac{1-\mu}{1-\mu(1-\lambda)}$, and it dominates bundling if $\chi \geq h_1(\lambda, \mu)$.

(ii) Case 1L dominates bundling if $\chi \geq h_{1\text{L}}(\lambda, \mu)$.

(iii) Case 2 is valid if $\chi \geq g_2(\lambda, \mu) = \frac{\beta\lambda\mu}{1+\beta\mu(1-\lambda)}$, and it dominates bundling if $\chi \geq h_2(\lambda, \mu)$.

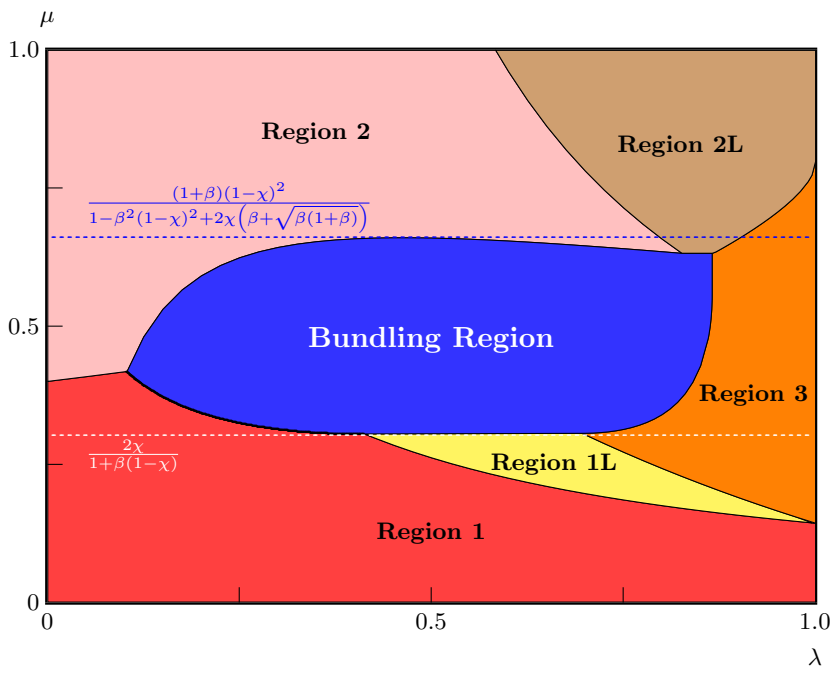

Figure 4: Equilibium Regions with Bundling ($\beta = 0.4$, $\chi = 0.2$)

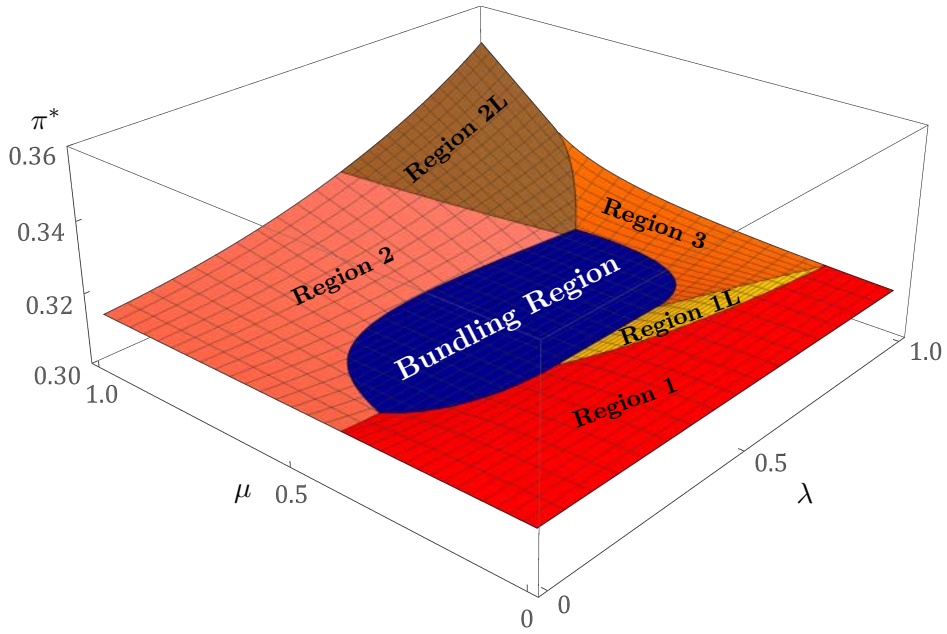

Figure 5: Vendor's Profit in Equilibrium with Bundling ($\beta = 0.4$, $\chi = 0.2$)

(iv) Case 2L dominates bundling if $\chi \geq h_{2L}(\lambda, \mu)$.

(v) Case 3 is valid if $\chi \leq g_3(\lambda,\mu) = \frac{(1+\beta)(\lambda(1+\mu(1+\beta(1-\mu)))-(1-\mu)(1+\beta\mu))}{2\lambda-\beta(1-\lambda)(1-\mu(1+\lambda-\beta(1-\mu(1-\lambda))))}$*, and it dominates bundling if* $\chi \geq h_3(\lambda,\mu)$.

In Lemma 2, the g-boundaries are case-appropriate validity conditions, whereas the h-boundaries are results of profit comparisons. Lemma 2 leads us to the following important result:

Theorem 1 *For moderate values of* λ *and* μ*, bundling can indeed be a viable strategy, as long as the marginal support cost,* χ*, is not too high. More specifically, bundling is optimal if and only if the following five conditions are all satisfied: (i)* $\chi < \max\{g_1, h_1\}$*, (ii)* $\chi < h_{1\mathrm{L}}$*, (iii)* $\chi < \max\{g_2, h_2\}$*, (iv)* $\chi < h_{2\mathrm{L}}$*, and (v)* $\chi \notin [h_3, g_3]$*, where* $g_i(\lambda,\mu)$*,* $i \in \{1,2,3\}$*, and* $h_j(\lambda,\mu)$*,* $j \in \{1,1\mathrm{L},2,2\mathrm{L},3\}$*, are all functions of* λ *and* μ *as defined in Lemma 2.*

Satisfying all the five conditions in Theorem 1 may seem like a tall order at first glance, but all of them can indeed be satisfied for a significant portion of the parameter space. What these conditions materially imply is simply this: as long as the service provision is sufficiently profitable—that is, as long as χ is sufficiently small—bundling could turn out to be an optimal strategy for the vendor. Essentially, the size of the bundling region (the blue lake in Figures 4 and 5) depends critically on χ. When χ is zero, the service provision is so profitable that bundling becomes optimal in the entire parameter space—the blue lake floods over the whole (λ,μ)-space in Figure 4. As χ increases, however, the blue lake starts shrinking, evaporating completely when χ becomes sufficiently large. This pattern becomes evident from the following corollary:

Corollary 2 *For bundling to be the vendor's optimal pricing strategy,* μ *must be moderate. More specifically, bundling is optimal only if* $\frac{2\chi}{1+\beta(1-\chi)} < \mu < \frac{(1+\beta)(1-\chi)^2}{1-\beta^2(1-\chi)^2+2\chi\left(\beta+\sqrt{\beta(1+\beta)}\right)}$*. Hence, bundling cannot be optimal if* $\chi > \bar{\chi}$*, where* $\bar{\chi}$ *is the smallest real root of:* $\frac{2\chi}{1+\beta(1-\chi)} = \frac{(1+\beta)(1-\chi)^2}{1-\beta^2(1-\chi)^2+2\chi\left(\beta+\sqrt{\beta(1+\beta)}\right)}$.

The two bounds of μ in Corollary 2 are also illustrated in Figure 4. It is easy to see that both the bounds are monotonic in χ, the lower bound monotonically increasing and the upper, monotonically decreasing, and the gap between them shrinking as χ increases. Therefore, the gap is the largest at $\chi = 0$, and it vanishes at $\chi = \bar{\chi}$, making bundling suboptimal beyond this point.

Theorem 1 and Corollary 2 provide several new insights. They tell us that bundling which was not optimal earlier (see Corollary 1) can now become an optimal pricing strategy for the vendor. Furthermore, bundling can be an appropriate tool for the vendor to put a floor on the immediate

profits and to thwart some of the losses brought on by the forum. They also tell us that, the higher the profit margin in provisioning support, the higher is the likelihood that such a hedging strategy could indeed work.

When $\chi < \bar{\chi}$, that is, when bundling could actually be optimal, the discrete jumps in the vendor's pricing strategy remain qualitatively the same throughout the parameter space. When the forum enters, the vendor initially sticks to its guns and ignores the entry—it continues to sell the software and its support separately. As the forum grows and gains in strength, at some point, the vendor's profit takes a serious hit. It is only then that the vendor changes course and starts selling the product and its support together as a bundle. The bundling strategy allows the vendor to weather the storm created by the substitutive effect of the forum. The vendor continues with this strategy as the forum keeps growing. However, a point is reached in this growth where the complementary effect of the forum takes control.[5] Now, the vendor changes course once again and starts selling separately. That, with the growth of the forum, a vendor's pricing strategy could swing back and forth between bundling and selling separately is a new insight and has not been discussed in the literature.

Finally, we turn our attention to public welfare. We find that the consumer surplus is not monotonic in λ or μ—there could be discrete jumps when the vendor's pricing strategy changes from one region to another. Ignoring these jumps, the following general pattern can be observed. Within a specific region, the consumer surplus is increasing in λ for low values of λ but is decreasing at higher values. When λ increases initially, more consumers tend to benefit from the forum, leading to an increase in consumer welfare. However, after λ increases to a substantial level, the vendor is able to extract a big portion of that benefit by pricing the product higher, which may lead to a decrease in the consumer surplus. In contrast, as μ increases, the benefit that the savvy segment derives from the forum increases. Hence, within a specific region, the consumer surplus increases with μ.

Despite the highly non-monotonic nature of the vendor's profit and consumer surplus, the total social surplus is somewhat better behaved in equilibrium. We illustrate it in Figure 6. For smaller values of μ—in Regions 1, 1L, and 3—the total surplus is monotonically increasing in both λ and

[5]The complementary effect can only be realized after a forum gains sufficient momentum. It is not necessary, however, that every forum survives, or thrives for that matter. A forum could as well die out or stagnate if it does not gather sufficient momentum; a complementary effect may not materialize in that case.

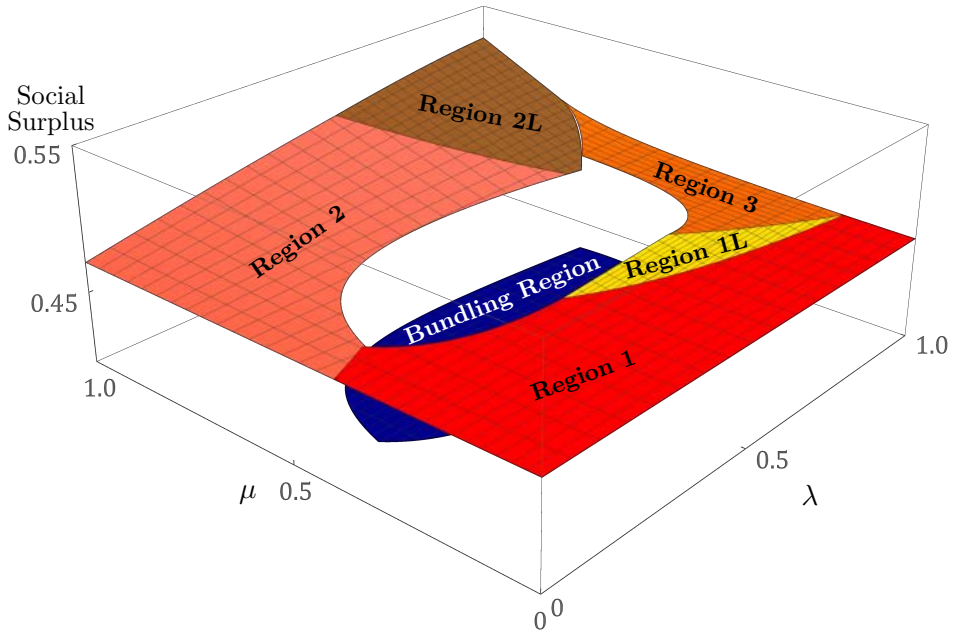

Figure 6: Total Social Surplus ($\beta = 0.4$, $\chi = 0.2$)

μ. For larger values of μ as well—in Regions 2 and 2L—the total surplus is increasing in μ; it is also increasing in λ as long as χ is not too low. The near-monotonicity is broken at the boundary between Regions 1 and 2, where there is a small dip in profit when moving from Regions 1 to 2. A similar dip is also observed at the upper boundary of Region 3.

In contrast to these minor discontinuities, a significant dip in social surplus occurs in the bundling region, that is, at moderate values of λ and μ. Why does welfare nosedive in the bundling region? In essence, when selling the product and its support as a bundle, the vendor turns away a significant number of consumers who would have purchased the product had it been sold separately. The resulting loss of surplus from the non-consumption of a zero marginal cost good is so large that it impacts the total surplus in a significant way. Put differently, even though the bundling strategy allows the vendor to put a floor on its own losses, the vendor does so only with a staggeringly negative repercussion on welfare.

Despite such immediate losses in total surplus, a strong forum turns out to be good for the society, eventually:

Theorem 2 *Social welfare is not monotonic in either λ or μ. It is at its lowest in the bundling*

region and at its maximum at $(\lambda, \mu) = (L, 1)$*, where:*

$$L = \begin{cases} 1, & \text{if } \chi \geq \frac{\beta(1+\beta)}{2(3-\beta)}, \\ \frac{(1+\beta)^2 - \sqrt{\frac{(1+\beta)(1+\beta-\chi)(1+\beta+3\chi)}{1+\chi(14-3\chi)+\beta(1+3\chi(2-\chi))}}}{\beta(2+\beta)}, & \text{otherwise.} \end{cases}$$

Theorem 2 is better visualized in Figure 7, where the (β, χ)-space is partitioned into two regions: (i) a green region where $L = 1$, and (ii) a red region where $L = \frac{(1+\beta)^2 - \sqrt{\frac{(1+\beta)(1+\beta-\chi)(1+\beta+3\chi)}{1+\chi(14-3\chi)+\beta(1+3\chi(2-\chi))}}}{\beta(2+\beta)} < 1$. There are a few interesting observations that can be made from Theorem 2 and Figure 7. First,

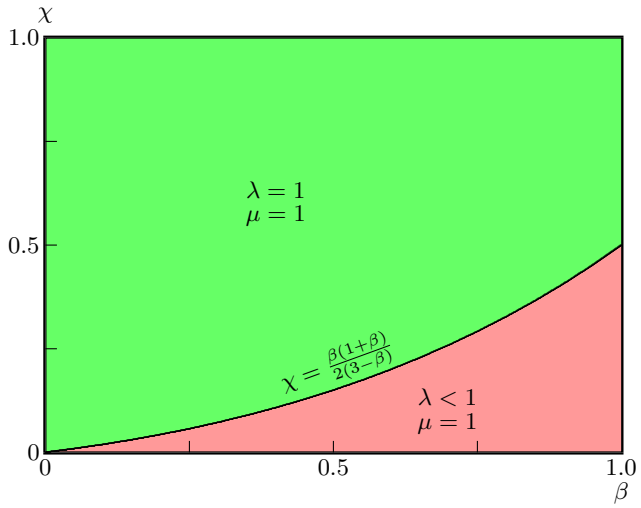

Figure 7: Partitions of the (β, χ) Space: Welfare-Maximizing λ and μ

irrespective of the value of β or χ, social welfare is maximized at $\mu = 1$. This makes sense. It is only when $\mu \to 1$ that the (free) forum finally gains its full effectiveness and becomes a close substitute for costly professional support. Naturally, when the forum becomes stronger, not only does it reward the vendor with a higher profit, but it also enhances the social surplus. Second, in the entire green region—that is, as long as χ is not too small and β, not too large—the social surplus is maximized when $\lambda = 1$. When the forum is highly effective, the vendor wants to offload as much of its costly support services as possible to the free forum; hence, it would prefer more people to become savvy, so that it does not have to incur the marginal cost of support but can still leverage the savvy consumers' extra WTP. This is why the vendor would prefer to see a large λ. It turns out that, as long as χ is not too small—more specifically, as long as $\chi \geq \frac{\beta(1+\beta)}{2(3-\beta)}$—it is also better for social welfare that a costly production process (professional support service) is

completely replaced by the free one (forum). In the red region, however, $\chi < \frac{\beta(1+\beta)}{2(3-\beta)}$; since support is relatively less expensive now, it is no longer socially optimal to replace support completely by the forum. Interestingly, even in this region, λ that maximizes social welfare is actually quite close to one, except when β and χ take on extreme values. The implication is clear. The existence of strong support forums is largely good for the economy, because they substitute for costly support services provided by the vendor.

At the same time, though, Theorem 2 tells us that the path towards a strong forum is not necessarily a smooth one. The overall non-monotonicy of the vendor's profit and social welfare—especially, the dramatic dip in social welfare in the bundling region—means that it is not just the vendor who faces the blues, but the society as a whole does so, too. However, the eventual outlook, provided that the forum survives the initial hiccups and gains sufficient momentum, appears to be quite good for both public welfare and private profits. Public policy about online forums must take into account this high level of non-monotonicy. Although support forums can eventually be good, there could indeed be some immediate pains; vendors, and even policymakers, may consider providing help and creating incentives for these forums to grow, and grow fast.

7 Possible Extensions

7.1 Penetration versus Effectiveness

So far, we have measured the forum's momentum or its strength using two parameters—penetration, λ, and effectiveness, μ—and have assumed that they are orthogonal. However, as mentioned in Section 5, they are likely to be highly dependent on each other—as more and more people join the forum, more questions are raised and resolved, making the forum more effective in resolving newer issues. Therefore, we can expect the forum's effectiveness, μ, to be proportional to the size of the savvy user segment, $\lambda\left(1-\frac{p}{1+\beta\mu}\right)$:

$$\mu(\lambda;\eta) = \eta\lambda\left(1-\frac{p}{1+\beta\mu(\lambda;\eta)}\right),$$

where $\eta > 0$, the network effect parameter, is small enough to ensure that $\mu(1;\eta) < 1$. The above expression can be solved to obtain:

$$\mu(\lambda;\eta) = \frac{\eta\lambda}{2} + \frac{\sqrt{1+\beta\eta\lambda(2-4p+\beta\eta\lambda)}-1}{2\beta} > 0, \tag{2}$$

which can now be substituted back into (1). The rest of the analysis proceeds in a manner similar to that in Sections 5 and 6, except that μ is now endogenized as above. The new objective function can be maximized over ρ, p, and s to obtain the equilibrium solution. We can then substitute p^* back into (2) to get the actual relationship between μ and λ in equilibrium. We illustrate this relationship in Figure 8 for four different values of η.

Fortunately, our original analysis in Section 5, where λ and μ were assumed to be orthogonal, can still be used to analyze the current situation where they are actually related to each other. To that end, we simply superimpose the regions in Figures 4 onto Figure 8. Earlier, when λ and

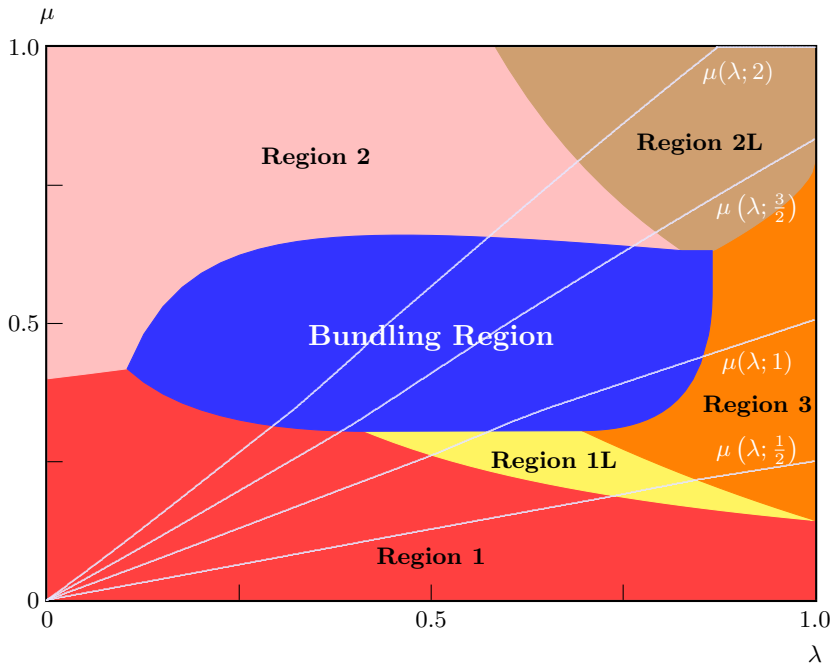

Figure 8: Equilibrium Path, $\mu(\lambda;\eta)$, for a Forum's Possible Growth ($\beta = 0.4$, $\chi = 0.2$)

μ were orthogonal, the entire (λ, μ)-space was relevant in equilibrium—depending on the state of the forum, the equilibrium could occur anywhere in the unit square in Figure 4. In contrast, now that λ and μ are related to each other by (2), the valid equilibrium cannot be just anywhere in

this space and must fall on the curve representing $\mu(\lambda;\eta)$, the equilibrium path along which the forum can possibly grow.[6] As long as this path passes through the bundling region (the blue lake in the figure), our earlier insights remain qualitatively valid: When the forum is very weak, the vendor chooses to sell the product and its support separately. However, when the forum's strength becomes moderate, the vendor changes course and adopts the bundling strategy. At even higher levels of growth, the vendor changes track once again and goes back to selling the product and its support separately. We find that when η is small, that is, when the network effect is minimal, the path may not pass through the blue lake. In that case, the forum never really gains enough momentum in terms of its effectiveness, and the vendor essentially continues to ignore its presence. Bundling is never a good strategy in that case. However, as the network effect picks up, the path crosses the bundling region, and the vendor ends up switching its pricing strategy back and forth along this path.

As before, whether bundling is a viable strategy depends critically on the marginal cost, χ. When χ goes down, the blue region expands, and the likelihood of $\mu(\lambda;\eta)$ passing through that region increases. On the other hand, when χ increases, the blue region shrinks, and along with it shrinks the possibility that $\mu(\lambda;\eta)$ passes through it. When χ is very large, the blue region disappears, and bundling is no longer a profitable strategy irrespective of the value of η. In summary, our main result that bundling can be a useful strategy, albeit only at moderate levels of growth of the forum, continues to hold even when penetration and effectiveness are no longer orthogonal.

7.2 Vendor's Investment and Support Quality

So far in our analysis, the quality level of vendor-provided support is considered fixed and normalized to one. However, it is entirely possible, especially in the face of shadow competition from the forum, that the vendor might invest in improving the quality of the support it provides. This it can do, for instance, by investing more in better infrastructure and technology for its support centers and providing better training to its support staff. We assume the following:

Assumption 6 *The vendor has the option of enhancing the quality of its support from* 1 *to* $1+\theta$ *by investing* $\frac{k\theta^2}{2}$.

[6] The restriction that the equilibrium must fall on the $\mu(\lambda;\eta)$ curve may make some of the five regions in Proposition 3 impossible. For example, in Figure 8, when $\eta \leq 1$, Regions 2 and 2L can no longer occur. Similarly, when $\eta \geq \frac{3}{2}$, Region 1L is not possible.

Notwithstanding Assumption 6, our analysis remains quite similar to that in Sections 5 and 6. To get the profit function for this setting, we simply need to replace β, μ, and χ in (1) with $\beta' = \beta(1+\theta)$, $\mu' = \frac{\mu}{1+\theta}$, and $\chi' = \frac{\chi}{1+\theta}$, respectively, and subtract the investment cost of $\frac{k\theta^2}{2}$. Maximizing this revised profit over ρ, p, s, and θ, we find that, as before, the parameter space gets partitioned into different regions. These regions are illustrated in Figure 9, with and without bundling. It can be readily verified from this figure that our results remain qualitatively similar

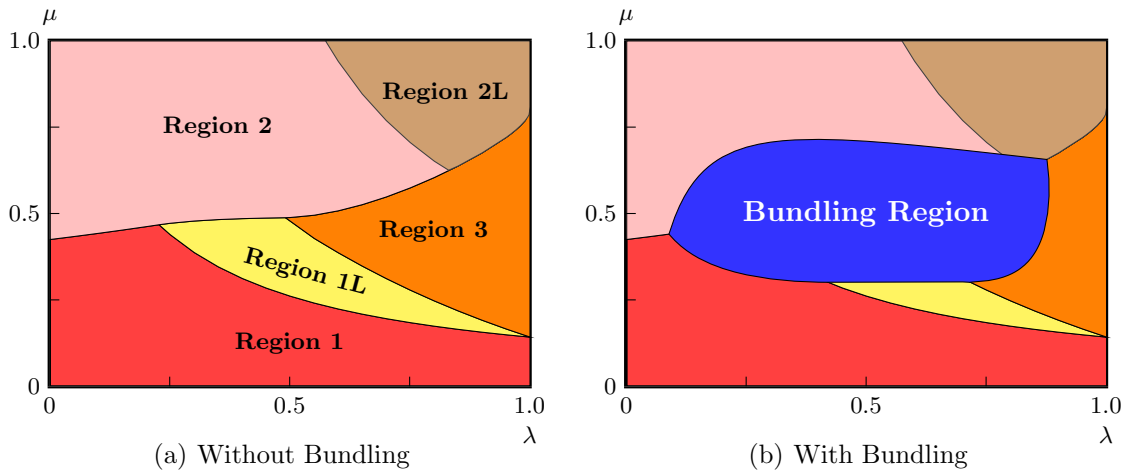

(a) Without Bundling (b) With Bundling

Figure 9: Equilibrium Regions with Investment in Support Quality ($\beta = 0.4$, $\chi = 0.2$, $k = 2$)

for this extension as well. When the forum enters, the vendor experiences a dip in its profit for moderate values of λ and μ. And, as before, this immediate loss can be addressed by the vendor if it chooses to sell the product and its support services simply as a bundle.

Interestingly, when the vendor has the option of investing extra into the quality level of its support services, it always chooses a $\theta^* > 0$, as long as k is not too large. The impact of k, the fixed cost parameter, on the vendor's adoption of bundling as a strategy is quite similar to the impact from marginal cost, χ. When k goes down, the vendor invests more in support quality. Since the vendor can recoup that investment quite effectively by bundling support with the product, bundling becomes more likely to be the dominant strategy at a lower k. As a result, the lake rises when k is small—just as it does when χ is small—and shrinks when k is larger. The only difference is that, unlike χ, a very large k by itself cannot make the lake disappear completely. This is because, when $k \to \infty$, we are back to our original analysis in Section 6 where $\theta = 0$.

8 Conclusion

Online communities providing support for various software products have proliferated in recent years. Their rapid emergence has prompted many researchers to explore how such communities form and grow, especially in terms of the incentives of users to join and contribute. However, the other side of the story—how a software vendor should change its pricing strategy in the face of this proliferation—has not received its due attention. This work is among the first to take up this elusive question. We consider a setting with two consumer types, a non-savvy group that depends critically on vendor-provided support and a savvy segment that can use these forums to substitute for the same. Intuitively, we expect this substitution to depress the demand for vendor-provided support, leading to a lower profit. However, as it turns out, this intuition is too simplistic in our context.

As revealed by our analyses, forums indeed play the role of a substitute, but they do so only to an extent. Interestingly, forums also play a complementary role—just as they substitute for the support provided by the vendor, they also act as a complement to the product itself. Thus, the implications for a vendor's pricing strategy depends critically on which role actually dominates. In the initial period of growth when the forum is not very popular or effective, its role as a substitute is more prominent, and its presence adversely affects the profit. As the forum becomes stronger in terms of popularity and effectiveness, this prominence only grows, taking a further toll on the profit. Only after a critical threshold is crossed, does the complementary role take over, and the profit starts to rise. In fact, when the forum is sufficiently strong, the vendor can make more money than it can in the absence of the forum, suggesting that the vendor has much to gain from a strong forum just as it has something to lose from a weaker one. This insight is novel, and it provides clear economic justifications in favor of interventions by a vendor towards strengthening existing consumer support communities and starting its own.

What explains this non-monotonicity? Essentially, when the forum is of low strength, charging a high price for the after-sales support is difficult because that would make the savvy group forgo support services. Moreover, as the forum gains strength, these consumers become even less willing to buy support as they are now happier to substitute the forum for vendor-provided support. Beyond a threshold, however, the vendor is happy to let them go, because it can then simply price

the product high to extract much of the surplus they are able to get from the forum. In other words, the complementary effect sets in and starts working in the opposite direction of the substitution effect.

As discussed above, the vendor's profit first decreases with the strength of the forum and then increases only after a threshold is crossed. Put another way, a forum that is of moderate strength can be highly detrimental to the vendor's profit. Fortunately, as we explain, the vendor can mitigate this detriment by bundling after-sales support with the product. This is indeed interesting. When the forum is very weak or strong, the vendor is better off not bundling the two, although exactly the opposite is preferable when the forum is of moderate strength. This unexpected finding has obvious implications for a software producer's pricing strategy.

The economics behind this phenomenon is also quite interesting. When the forum is of moderate strength, the vendor finds itself in a bind—it is neither able to charge a high price for the product nor is it able to do so for support. In essence, when the forum is of moderate strength, the savvy group, which has the required ability to enjoy the free benefits offered by the forum, values the product quite a bit more than the non-savvy. At the same time, however, for exactly the same reason, the savvy is now somewhat less willing to pay for vendor-provided support. The net effect is an infusion of a valuation heterogeneity that works against the vendor. Logically, bundling ought to be effective in such a situation, as it can mitigate this heterogeneity between the two consumer segments—unlike the individual valuations for the product and support, the valuation for the bundle does not depend on a consumer's savvy. This is precisely what makes bundling useful. Of course, bundling has its own limitations. Notably, the vendor has little to gain from bundling at either extreme. On one hand, when the forum is absent, or its presence minimal, the marginal cost of support makes it a natural candidate for being sold separately at a premium. On the other, when the forum is strong, bundling drives even savvy consumers towards the costly vendor-provided support, thereby failing to leverage on the benefits of a strong forum. These insights are clearly borne out by our analysis. They also remain valid when we incorporate network effects or allow the vendor another lever in the form of tweaking support quality.

We find that, eventually, a strong forum is not just beneficial to the vendor but is also so from the broader perspective of public welfare. This is true despite the fact that the growth path of the forum is not necessarily smooth, and there could be some immediate loss in welfare. Put differently,

both the vendor and policymakers should welcome the formation and growth of support forums, but they must also have a strategy to deal with the proximate blues.

Of course, forums not only help savvy users with necessary product-related support, but they also provide several other social benefits. Since a forum facilitates brainstorming and active participation, it can often be a breeding ground for innovative solutions and workarounds towards difficult issues related to a product. Specifically, in situations where the vendor-provided support services do not have a ready-made solution for a newly identified issue with the product, a forum that brings together many knowledgeable people may sometimes be able to converge to an appropriate solution even before the vendor can. Furthermore, based on user feedback and comments it generates and stores, a forum may hold the keys to a better product design in the future. Forums can also help find newer ways of utilizing existing features within a product and discover product glitches previously unknown to everyone. Huang et al. (2018, p.214) summarize this well: "These communities enable the rapid diffusion of best practices, strengthen customer relationships, gather inputs and feedback for new product and service developments, reduce support costs, and incubate user-driven innovations." Admittedly, we do not consider such positive influences of the forum within our setup and only underestimate the overall value a forum creates. The inclusion of such effects would, therefore, bias our results along a predictable direction.

In sum, the economics of support forums is a new topic, and our work seeks to contribute to this emerging domain. Pricing and bundling should not be the only consideration, however. There are several other issues, including versioning, piracy, and competition, which all can play a role in how a forum affects the vendor or the society. One could also look into the vendor's incentive towards investing in a better quality product. Our model does not incorporate these issues to keep potential confounding factors at bay, so our results should be interpreted accordingly. In future work, we will seek to address some these issues, to develop a more comprehensive picture of the context.

References

Adam, W.J., J.L. Yellen. 1976. Commodity bundling and the burden of monopoly. *Quarterly Journal of Economics* **90**(3) 475–498.

Bakos, Y., E. Brynjolfsson. 2000. Bundling and competition on the internet. *Marketing Science* **19**(1) 63–82.

Bhargava, H.K., V. Choudhary. 2008. Research note—When is versioning optimal for information goods? *Management Science* **54**(5) 1029–1035.

Dewan, R., M.L. Freimer. 2003. Consumers prefer bundled add-ins. *Journal of Management Information Systems* **20**(2) 99–111.

Huang, P., A. Tafti, S. Mithas. 2018. Platform sponsor investments and user contributions: The role of knowledge seeding. *MIS Quarterly* **42**(1) 213–240.

Huang, P., Z. Zhang. 2016. Participation in open knowledge communities and job-hopping: Evidence from enterprise software. *MIS Quarterly* **40**(3) 785–806.

Jabr, W., R. Mookerjee, Y. Tan, V.S. Mookerjee. 2014. Leveraging philanthropic behavior for customer support: The case of user support forums. *MIS Quarterly* **38**(1) 187–208.

Khansa, L., X. Ma, D. Liginlal, S.S. Kim. 2015. Understanding members' active participation in online question-and-answer communities: A theory and empirical analysis. *Journal of Management Information Systems* **32**(2) 162–203.

Lappas, T., C. Dellarocas, S.N. Derakhshan. 2017. Reputation and contribution in online question-answering communities. Social Science Research Network. URL `https://ssrn.com/abstract=2918913`.

Lu, Y., P. V. Singh, B. Sun. 2017. Is a core-periphery network good for knowledge sharing? A structural model of endogenous network formation on a crowdsourced customer support forum. *MIS Quarterly* **31**(2) 607–628.

Pang, M.-S., H. Etzion. 2012. Research note—Analyzing pricing strategies for online services with network effects. *Information Systems Research* **23**(4) 1364–1377.

Parker, G.G., M.W. Van Alstyne. 2005. Two-sided network effects: A theory of information product design. *Management Science* **51**(10) 1494–1504.

Zhao, L., B. Detlor, C.E. Connelly. 2016. Sharing knowledge in social q&a sites: The unintended consequences of extrinsic motivation. *Journal of Management Information Systems* **33**(1) 70–100.

Appendix: Proofs

Proof of Proposition 1

Since $\pi = p(1-p) + (s-c)\left(1 - \frac{s}{\beta}\right)$, the Hessian matrix is given by:

$$H = \begin{bmatrix} -2 & 0 \\ 0 & -\frac{2}{\beta} \end{bmatrix},$$

which has a positive determinant, and its first term is negative. In other words, the profit function is concave in (p, s), and (p^*, s^*) can be easily found from the following two first order conditions.

$$\frac{\partial \pi}{\partial p} = 1 - 2p = 0 \quad \text{and} \quad \frac{\partial \pi}{\partial s} = 1 + \frac{c}{\beta} - \frac{2s}{\beta} = 0 \quad \Leftrightarrow \quad (p^*, s^*) = \left(\frac{1}{2}, \frac{\beta(1+\chi)}{2}\right),$$

where $\chi = \frac{c}{\beta}$. Substituting (p^*, s^*) into π provides the desired $\pi^* = \frac{1+\beta(1-\chi)^2}{4}$. ■

Proof of Proposition 2

The vendor's profit $\pi_b = (p_b - c)\left(1 - \frac{p_b}{1+\beta}\right)$ is concave in p because $\frac{\partial^2 \pi_b}{\partial p_b^2} = -\frac{2}{1+\beta} < 0$, for all $\beta \in (0,1)$. Therefore, p_b^* can be obtained by solving the first order condition:

$$\frac{d\pi_b}{dp_b} = 1 - \frac{2p_b}{1+\beta} + \frac{c}{1+\beta} = 0 \quad \Leftrightarrow \quad p_b^* = \frac{1+\beta(1+\chi)}{2},$$

where, as before, $\chi = \frac{c}{\beta}$. Substituting p_b^* in π_b provides $\pi_b^* = \frac{(1+\beta(1-\chi))^2}{4(1+\beta)}$. ■

Proof of Corollary 1

It follows directly from Propositions 2 and 1, by comparing the profit from bundling (in Proposition 2) with that from selling the product and support separately (in Proposition 1). ■

Proof of Lemma 1

In Case 4, $p > \frac{s}{\beta}$ and either $\frac{p}{1+\beta\mu} \geq \frac{s}{\beta(1-\mu)}$ or $\frac{s}{\beta(1-\mu)} \geq 1$.

When $p > \frac{s}{\beta}$, if the equilibrium solution were to satisfy $\frac{s}{\beta(1-\mu)} \geq 1$, no savvy consumer would consider purchasing support but all non-savvy consumers would. Therefore, a slight increase in s would not affect the total demand for support, as long as $p > \frac{s}{\beta}$ holds. In other words, the vendor would be able to increase its profit by slightly increasing s, implying that the solution could not have been an equilibrium to begin with. So, $\frac{s}{\beta(1-\mu)} \geq 1$ is not possible.

Now, if the equilibrium were to satisfy $p > \frac{s}{\beta}$ and $\frac{p}{1+\beta\mu} \geq \frac{s}{\beta(1-\mu)}$, together they would imply that all consumers who purchase the product would end up buying the professional support if the need for it arises after using the product. The resulting profit to the vendor in Case 4 can the be written as:

$$\pi_4 = (p + (1-\rho)(s-c))\left(\lambda\left(1 - \frac{p}{1+\beta\mu}\right) + (1-\lambda)(1-p)\right) - \kappa(\rho).$$

It is easy to see that π_4 is increasing in s, so no such interior solution is possible in this case. ■

Proof of Proposition 3

- Region 1: The first order conditions in this region are:

$$\frac{\partial \pi_1}{\partial p} = 1 - \frac{2p(1+\beta\mu(1-\lambda))}{1+\beta\mu} = 0 \quad \text{and} \quad \frac{\partial \pi_1}{\partial s} = 1 + \frac{1-\mu(1-\lambda)}{1-\mu}\left(\chi - \frac{2s}{\beta}\right) = 0.$$

  Solving them for p and s we get the equilibrium prices for Region 1. The second order conditions have also been verified.

- Region 1L: We set $s = \beta p$ and substitute that in π_1. Then, the first order condition becomes:

$$\frac{d\pi_1}{dp} = 1 + \beta\left(1 + \frac{\chi(1-\mu(1-\lambda))}{1-\mu}\right) - 2p\left(1 + \beta\left(1 + \frac{\lambda\mu(\mu(1+\beta))}{(1-\mu)(1+\beta\mu)}\right)\right) = 0.$$

  Since $\frac{d^2\pi_1}{dp^2}$ is clearly negative, $p_{1\mathrm{L}}^*$ can be determined by solving the first order condition for p. Finally, $s_{1\mathrm{L}}^* = \beta p_{1\mathrm{L}}^*$.

- Region 2: In this regions, the first order conditions are:

$$\frac{\partial \pi_2}{\partial p} = 1 - \frac{2p(1+\beta\mu(1-\lambda))}{1+\beta\mu} \quad \text{and} \quad \frac{\partial \pi_2}{\partial s} = (1-\lambda)\left(1+\chi-\frac{2s}{\beta}\right).$$

  By simultaneously solving them for p and s, we get the equilibrium prices for Region 2. Second order conditions were also verified.

- Region 2L: We set $s = \beta p$ and substitute that in π_2. Then, the first order condition becomes:

$$\frac{d\pi_2}{dp} = 1 + \beta(1+\chi)(1-\lambda) - \frac{2p(1+\beta(1-\lambda)(1+\mu(1+\beta)))}{1+\beta\mu} = 0.$$

  Since $\frac{d^2\pi_2}{dp^2}$ is clearly negative, $p^*_{2\mathrm{L}}$ can be determined by solving the first order condition for p. Finally, $s^*_{2\mathrm{L}} = \beta p^*_{2\mathrm{L}}$.

- Region 3: The first order conditions in this region are:

$$\frac{\partial \pi_3}{\partial p} = 1+\beta(1-\lambda)\left(\chi-\frac{s}{\beta}\right)-\frac{2p(1+\beta\mu(1-\lambda))}{1+\beta\mu} = 0 \quad \text{and} \quad \frac{\partial \pi_3}{\partial s} = 1+\frac{\lambda\chi}{1-\mu}-p(1-\lambda)-\frac{2s\lambda}{\beta(1-\mu)} = 0.$$

  By simultaneously solving them for p and s, we get the equilibrium prices for Region 3. Second order conditions were also verified. ■

Proof of Lemma 2

(i) Case 1 is valid if $p^*_1 \le \frac{s^*_1}{\beta}$ and $\frac{s^*_1}{\beta(1-\mu)} < 1$. Substituting p^*_1 and s^*_1 from Proposition 3 yields that $p^*_1 \le \frac{s^*_1}{\beta}$ is equivalent to $\chi \ge g_1(\lambda,\mu)$, where $g_1(\lambda,\mu) = \frac{1+\beta\mu}{1+\beta\mu(1-\lambda)} - \frac{1-\mu}{1-\mu(1-\lambda)}$. We now show that the other validity condition, namely $\frac{s^*_1}{\beta(1-\mu)} < 1$, is not relevant in our context. To see this clearly, suppose that the optimal solution is in Case 1 but (p^*_1, s^*_1) violates the condition $\frac{s^*_1}{\beta(1-\mu)} < 1$. Then, from (1), we see that:

$$\pi_1(p^*_1,s^*_1) = \pi_2(p^*_1,s^*_1) + \lambda(s^*_1-c)\left(1-\frac{s^*_1}{\beta(1-\mu)}\right) \le \pi_2(p^*_1,s^*_1) \le \pi_2(p^*_2,s^*_2).$$

Evidently, Case 2 provides a better profit, which contradicts our earlier supposition that Case 1 is optimal.

When Case 1 is valid, we can compare the profit with the bundling profit; algebraically, $\pi^*_1 \ge \pi^*_b$ is equivalent to $\chi \ge h_1(\lambda,\mu)$, where

$$h_1(\lambda,\mu) = \sqrt{\frac{\lambda\mu^2(1-\lambda)(1-\mu)(1+\beta)^2}{(1-\mu(1-\lambda))(1+\beta\mu(1-\lambda))(1-\mu(1-\lambda(1+\beta)))}}.$$

(ii) Case 1L is valid if $\frac{s^*_{1L}}{\beta(1-\mu)} < 1$, which we can again show to be irrelevant using an argument analogous to the one used above. When comparing the profits, we find that $\pi^*_{1L} \ge \pi^*_b$ is algebraically equivalent to $\chi \ge h_{1L}(\lambda,\mu)$, where

$$h_{1L}(\lambda,\mu) = \frac{(1+\beta)\left(\sqrt{(1-\mu)(1+\beta\mu)(1-\mu(1-\beta+\beta\mu(1-\lambda)))}-(1-\mu)\right)}{\beta(2-\mu(2-\beta-\lambda+\beta\mu(1-\lambda)))}.$$

(iii) Case 2 is valid if $p_2^* \leq \frac{s_2^*}{\beta}$, since the condition $\frac{s_2^*}{\beta(1-\mu)} \geq 1$ can be shown to be irrelevant, as before. Substituting p_2^* and s_2^* from Proposition 3 yields $\chi \geq g_2(\lambda,\mu) = \frac{\beta\lambda\mu}{1+\beta\mu(1-\lambda)}$. When Case 2 is valid, we can compare the profit with the bundling profit; algebraically, $\pi_2^* \geq \pi_b^*$ is equivalent to $\chi \geq h_2(\lambda,\mu)$, where

$$h_2(\lambda,\mu) = \frac{\sqrt{\frac{\lambda(1+\beta)(1-\mu(1-\beta-\lambda))}{1+\beta\mu(1-\lambda)}} - \lambda(1+\beta)}{1-\lambda(1+\beta)}.$$

(iv) Case 2L is valid if $\frac{s_{2L}^*}{\beta(1-\mu)} \geq 1$, a condition which is again irrelevant. Furthermore, $\pi_{2L}^* \geq \pi_b^*$ is algebraically equivalent to $\chi \geq h_{2L}(\lambda,\mu)$, where

$$h_{2L}(\lambda,\mu) = \frac{1+\beta(2-\lambda+\beta(1-\lambda)+\mu(1+\beta)(2+\beta)(1-\lambda)) - \sqrt{(1+\beta)(1+\beta\mu(5-4\lambda+4\beta(1-\lambda)))(1+\beta(1-\lambda)(1+\mu(1+\beta)))}}{\beta(2-\lambda+\beta(1-\lambda)(1+\mu(1+\beta)))}.$$

(v) Case 3 is valid if $p_3^* > \frac{s_3^*}{\beta}$, since $\frac{s_3^*}{\beta(1-\mu)} < 1$ is once again not relevant. The condition $p_3^* > \frac{s_3^*}{\beta}$ is equivalent to $\chi < g_3(\lambda,\mu)$, where $g_3(\lambda,\mu) = \frac{(1+\beta)(\lambda(1+\mu(1+\beta(1-\mu)))-(1-\mu)(1+\beta\mu))}{2\lambda-\beta(1-\lambda)(1-\mu(1+\lambda-\beta(1-\mu(1-\lambda))))}$. Finally, $\pi_3^* \geq \pi_b^*$ is algebraically equivalent to $\chi \geq h_3(\lambda,\mu)$, where

$$h_3(\lambda,\mu) = \frac{(1+\beta)(1-\mu)\left(\beta^2\mu(1-\lambda)^2(1-\mu)(1-2\xi)-2\lambda(1-\lambda-4\xi)+\beta(1-\lambda)((1-\mu)(1-2\xi)-\lambda(1+\mu-2\xi(1+3\mu)))\right)}{4\lambda^2+\beta^3\mu(1-\lambda)^2(1-\mu)^2+\beta^2(1-\lambda)((1-\mu)^2+4\lambda^2\mu-\lambda(1-\mu)(1+3\mu))-4\beta\lambda(1-\mu-\lambda(1+\mu(1-\lambda)))}, \text{ and}$$
$$\xi = \sqrt{\frac{\lambda\mu(1-\lambda)(\lambda-\beta(1-\lambda-\mu))}{(1-\mu)(4\lambda-\beta^2\mu(1-\mu)(1-\lambda)^2-\beta(1-\lambda)(1-\lambda-\mu-3\lambda\mu))}}.$$

■

Proof of Theorem 1

Bundling would be optimal if the bundling profit is higher than all the five cases of selling the product and its support separately. We consider each case separately:

(i) From Lemma 2, we know that bundling dominates Case 1 if either Case 1 is not valid or if the bundling profit is higher, that is, if either $\chi < g_1$ or $\chi < h_1$, which is equivalent to $\chi < \max\{g_1, h_1\}$.

(ii) From Lemma 2, bundling dominates Case 1L if $\chi < h_{1\text{L}}$.

(iii) Again, from Lemma 2, bundling dominates Case 2 if either Case 2 is not valid or if the bundling profit is higher, that is, if either $\chi < g_2$ or $\chi < h_2$, which is equivalent to $\chi < \max\{g_2, h_2\}$.

(iv) From Lemma 2, bundling dominates Case 2L if $\chi < h_{2\text{L}}$.

(v) Finally, bundling is better than Case 3 if either $\chi \geq g_3$ or $\chi \leq h_3$, that is, if $\chi \notin [h_3, g_3]$.

When all of the above five conditions are satisfied at the same time, bundling dominates all the cases of selling separately, making bundling the optimal choice in that case. ■

Proof of Corollary 2

The upper boundary of the bundling region is characterized by the function separating Region 2 from the bundling region; in the proof of Lemma 2, this boundary was found to be $\chi = h_2(\lambda,\mu)$, which can be solved for μ. This way, the boundary can also be expressed as $\mu = H_2(\lambda)$, where:

$$H_2(\lambda) = \frac{\lambda(1+\beta)(1-\chi)^2 - \chi^2}{\lambda(1+\beta)(1-\beta(1-\lambda)) + 2\beta\chi\lambda(1+\beta)(1-\lambda) + \beta\chi^2(1-\lambda)(1-\lambda(1+\beta))}.$$

To find the maximum of $H_2(\lambda)$, we solve $\frac{\partial H_2}{\partial \lambda} = 0$, which has two roots, only one of which, $\lambda = \frac{\chi\left(\beta\chi+\sqrt{\beta(1+\beta)}\right)}{\beta(1+\beta)(1-\chi)^2}$, is the maxima, because, for all $\beta, \chi \in (0,1)$:

$$\left.\frac{\partial^2 H_2}{\partial \lambda^2}\right|_{\lambda=\frac{\chi\left(\beta\chi+\sqrt{\beta(1+\beta)}\right)}{\beta(1+\beta)(1-\chi)^2}} = -\frac{2\beta^3(1+\beta)^4(1-\chi)^6\left(1-\beta^2(1-\chi)^2+2\chi\left(\beta+\sqrt{\beta(1+\beta)}\right)\right)}{\chi\left(\sqrt{\beta(1+\beta)}\left(1-\beta^2(1-\chi)^2\right)+2\beta\chi\left(1+\beta+\sqrt{\beta(1+\beta)}\right)\right)^3} < 0.$$

Therefore, the upper bound of μ in the bundling region can be expressed as:

$$\bar{\mu} = H_2\left(\frac{\chi\left(\beta\chi+\sqrt{\beta(1+\beta)}\right)}{\beta(1+\beta)(1-\chi)^2}\right) = \frac{(1+\beta)(1-\chi)^2}{1-\beta^2(1-\chi)^2+2\chi\left(\beta+\sqrt{\beta(1+\beta)}\right)}.$$

Now, the lower bound of bundling region is characterized by the function separating the bundling region from Region 1L; according to Lemma 2, this boundary is given by $\chi = h_{1\mathrm{L}}(\lambda, \mu)$, which can be solved for μ, so the boundary can also be expressed as $\lambda = H_{1\mathrm{L}}(\mu)$, where:

$$H_{1\mathrm{L}}(\mu) = \frac{(1-\mu)(1+\beta(1+\chi))(\mu(1+\beta(1-\chi))-2\chi)}{\beta\mu\chi^2(1+\beta\mu)}.$$

Bundling is more profitable than Case 1L if λ is below this threshold. We first take a derivative of this threshold with respect to μ:

$$H'_{1\mathrm{L}}(\mu) = \frac{(1+\beta+\beta\chi)\left(\chi(2+\beta\mu(4-\mu(1-\beta)))-\mu^2(1+\beta)^2\right)}{\beta\chi^2\mu^2(1+\beta\mu)^2},$$

which is positive if $\chi > \hat{\chi} = \frac{\mu^2(1+\beta)^2}{2+\beta\mu(4-\mu(1-\beta))}$. Now, $\frac{\partial\hat{\chi}}{\partial\mu} = \frac{4\mu(1+\beta\mu)(1+\beta)^2}{(2+\beta\mu(4-\mu(1-\beta)))^2} > 0$, for all $\beta, \mu \in (0,1)$. So, for a given β and χ, $(\chi - \hat{\chi})$ can change sign at most once as μ increases. Next, it is easy to see that the numerator of $H_{1\mathrm{L}}(\mu)$ becomes negative and the denominator, zero, at $\mu = 0$, implying that $H_{1\mathrm{L}}(\mu) \to -\infty$ as $\mu \to 0$. Therefore, $H_{1\mathrm{L}}(\mu)$ crosses the μ-axis at most once. Since λ cannot be negative, the minimum possible value of μ can be found by setting $H_{1\mathrm{L}}(\mu) = 0$ and solving for μ; this lower limit is given by:

$$\underline{\mu} = \frac{2\chi}{1+\beta(1-\chi)}.$$

Note that $\underline{\mu} < \mu < \bar{\mu}$ provides a necessary, but not a sufficient, condition for bundling—bundling can be optimal *only if* $\mu \in (\underline{\mu}, \bar{\mu})$. Of course, if $\underline{\mu} > \bar{\mu}$, the interval $(\underline{\mu}, \bar{\mu})$ would no longer exist and the necessary condition for bundling would certainly be violated. Therefore, by equating $\underline{\mu} = \bar{\mu}$ and solving for χ, we can find the threshold on χ beyond which bundling cannot be optimal. Therefore, to complete the proof, we must show that $\underline{\mu} = \bar{\mu}$ has a unique solution in our context. First, it can be verified that $\underline{\mu} = \bar{\mu}$ is a cubic equation in χ, with a discriminant of

$$D = 4(1+\beta)^3\left(3-20\sqrt{\beta(1+\beta)}-\beta\left(31+4\sqrt{\beta(1+\beta)}-8\beta\left(3+7\sqrt{\beta(1+\beta)}+\beta\left(9+4\beta+4\sqrt{\beta(1+\beta)}\right)\right)\right)\right).$$

When $0.015591 < \beta < 0.49145$, D is negative, implying that only one root is real and the other two are imaginary. When $\beta \leq 0.015591$, $D \geq 0$ and there are three real roots, of which only one is in $(0,1)$ and the other two roots are both greater than 1. At the other end, when $\beta \geq 0.49145$, D is positive and, again, there are three real roots; of them, only one is in $(0,1)$ and the other two roots are both less than 0. Therefore, the smallest positive root is the unique solution in our context. ■

Proof of Theorem 2

Let $\phi_1 = \max\left\{p, \min\left\{1, \frac{s}{\beta}\right\}\right\}$ and $\phi_2 = \max\left\{\frac{p}{1+\beta\mu}, \min\left\{1, \frac{s}{\beta(1-\mu)}\right\}\right\}$. Then, the social surplus can be obtained from:

$$\sigma = (1-\lambda)\int_{p}^{1} v\,dv \;+\; \lambda\int_{\frac{p}{1+\beta\mu}}^{1} v(1+\beta\mu)\,dv \;+\; (1-\lambda)\int_{\phi_1}^{1}\beta(v-\chi)\,dv \;+\; \lambda\int_{\phi_2}^{1}\beta(v(1-\mu)-\chi)\,dv.$$

After substituting the appropriate values of p and s from Proposition 3, we get:

$$\sigma = \begin{cases} \sigma_1 = \frac{3\beta\chi((1-\mu)(\chi-2)+\chi\lambda\mu)}{8(1-\mu)} + \frac{(1+\beta)(3-\mu(1-\lambda)(3-\beta(3-\mu(3-4\lambda))))}{8(1-\mu(1-\lambda))(1+\beta\mu(1-\lambda))}, & \text{Case 1,} \\ \sigma_{1\mathrm{L}} = \frac{3(1+\beta\mu)(\beta\chi(1-\mu(1-\lambda)))^2}{8(1+\beta)(1-\mu)(1-\mu(1-\beta(1-\mu(1-\lambda))))} - \frac{2\beta\chi(3-\mu(3+\lambda-3\beta(1-\mu(1-\lambda))))}{8(1-\mu(1-\beta(1-\mu(1-\lambda))))} + \\ \qquad \frac{(1+\beta)(3-\mu(3-\beta(3(1-\mu)+4\lambda\mu)))}{8(1-\mu(1-\beta(1-\mu(1-\lambda))))}, & \text{Case 1L,} \\ \sigma_2 = \frac{3\beta\chi(1-\lambda)(\chi-2)}{8} + \frac{3(1+\beta(1-\lambda+\mu))+\beta^2\mu(1-\lambda)(3-3\lambda+4\lambda\mu)}{8(1+\beta\mu(1-\lambda))}, & \text{Case 2,} \\ \sigma_{2\mathrm{L}} = \frac{3(1+\beta\mu)(1+\beta(1-\lambda)(1-\chi))^2+4\beta^2\lambda\mu(1-\lambda)(2\chi+\mu(1+\beta))}{8(1+\beta(1-\lambda)(1+\mu(1+\beta)))}, & \text{Case 2L,} \\ \sigma_3 = \frac{\lambda\left(a_0+a_1\beta+a_2\beta^2+a_3\beta^3+a_4\beta^4+a_5\beta^5\right)}{2(1-\mu)(4\lambda-\beta^2\mu(1-\mu)(1-\lambda)^2-\beta(1-\lambda)(1-\lambda-\mu-3\lambda\mu))^2}, & \text{Case 3,} \end{cases}$$

where

$$\begin{aligned} a_0 &= 12\lambda(1-\mu), \\ a_1 &= (1-\mu)\left(\mu(1-\lambda)(5+19\lambda) - 5 + 22\lambda - 5\lambda^2\right) - 8\chi\lambda(2+\lambda)(1-\mu) + 12\chi^2\lambda^2, \\ a_2 &= (1-\lambda)\left((1-\mu)\left(\lambda(7+\mu(32-3\mu)) - 6 - 2\mu + 8\mu^2 - \lambda^2(1+3\mu)^2\right) + \right. \\ &\qquad \left. 2\chi(1-\mu)(3(1-\mu) - \lambda(3+\mu(15+2\lambda))) - \lambda\chi^2(1-\lambda-\mu-23\lambda\mu)\right), \\ a_3 &= (1-\lambda)\left((1-\mu)\left(\lambda\left(2+\mu\left(17+\mu-8\mu^2\right)\right) - (1-\mu)(1+2\mu(5+\mu)) - \right.\right. \\ &\qquad \left. \lambda^2(1+3\mu)(1+\mu(5-2\mu))\right) - \chi^2(1-\lambda)\left(1-\lambda-\mu(2-\lambda)(1-2\lambda) + \right. \\ &\qquad \left.\left. \mu^2(1-2\lambda(2+5\lambda))\right) - 2\chi(1-\lambda)(1-\mu)(\lambda-(1-\mu)(1+6\mu)+\lambda\mu(4+7\mu))\right), \\ a_4 &= \mu(1-\mu)(1-\lambda)^2\left(2\chi\left(2+\mu-3\mu^2-\lambda(2+3\mu(1-\mu))\right) - \chi^2(1-\lambda)(2-\mu(2-3\lambda)) - \right. \\ &\qquad \left. (1-\mu)(2+\mu(3-\mu)-\lambda(2+\mu(7-\mu)))\right), \text{ and} \\ a_5 &= \mu^2(1-\mu)(1-\lambda)^3\left(\mu(2-\chi(2-\chi(1-\lambda))) - \mu^2 - (1-\chi)^2\right). \end{aligned}$$

Also, we consider the bundling region, where the social surplus can be written as:

$$\sigma_b = \int_{\frac{p_b}{1+\beta}}^{1} (v(1+\beta)-\chi)\,dv = \frac{(1+\beta(1-\chi))(3\beta(1+\beta)-\beta\chi(4-\beta))}{8\beta(1+\beta)}.$$

First, in order to show that social surplus is the lowest in the bundling region, we will show that $\sigma_b \le \sigma_i$, for all $i \in \{1, 1\mathrm{L}, 2, 2\mathrm{L}, 3\}$:

- Since $\sigma_1 \ge \sigma_1|_{\mu=0}$, we have

$$\sigma_1 - \sigma_b \ge \sigma_1|_{\mu=0} - \sigma_b = \frac{\chi\left(4\left(1-\beta^2(1-\chi)\right)-\beta\chi\right)}{8(1+\beta)} \ge 0.$$

- Because $\sigma_{1\mathrm{L}} \geq \sigma_{1\mathrm{L}}|_{\mu=0}$, we have

$$\sigma_{1\mathrm{L}} - \sigma_b \geq \sigma_{1\mathrm{L}}|_{\mu=0} - \sigma_b = \frac{\chi(1-\beta)(1+\beta(1-\chi))}{2(1+\beta)} \geq 0.$$

- In this case, $\sigma_2 \geq \sigma_2|_{\lambda=0}$, so we have

$$\sigma_2 - \sigma_b \geq \sigma_2|_{\lambda=0} - \sigma_b = \frac{\chi\left(4\left(1-\beta^2(1-\chi)\right)-\beta\chi\right)}{8(1+\beta)} \geq 0.$$

- Since $\sigma_{2\mathrm{L}} \geq \sigma_{2\mathrm{L}}|_{\lambda=0}$, we have

$$\sigma_{2\mathrm{L}} - \sigma_b \geq \sigma_{2\mathrm{L}}|_{\lambda=0} - \sigma_b = \frac{\chi(1-\beta)(1+\beta(1-\chi))}{2(1+\beta)} \geq 0.$$

- Since σ_3 is increasing in both λ and μ, it cannot be any lower than its value at the lower boundary of Region 3, implying σ_3 is no smaller than $\sigma_{1\mathrm{L}}$. Therefore, $\sigma_3 - \sigma_b \geq \sigma_{1\mathrm{L}} - \sigma_b \geq 0$.

Clearly then, social surplus is always lowest in the bundling region.

We will now show in two steps that social surplus is maximized at $(\lambda, \mu) = (L, 1)$. First, we will show that the social surplus at $(\lambda, \mu) = (1, 1)$, denoted $\hat{\sigma}$, is higher than that in Regions 1, 1L, 2, and 3, as well as the surplus in the bundling region. Next, we will show that the social surplus at $(\lambda, \mu) = (L, 1)$ cannot be any lower than $\hat{\sigma}$, implying that the highest surplus is realized at $(L, 1)$.

We start by noting that the point $(1, 1)$ belongs to Region 2L. At that point, since $\lambda = \mu = 1$, there are only forum-savvy consumers, and they do not purchase the vendor-provided support any more. Instead, they buy the software product at a price of $\hat{p} = \frac{1+\beta}{2}$. The resulting social surplus at $(\lambda, \mu) = (1, 1)$ can be found to be:

$$\hat{\sigma} = \int_{\frac{\hat{p}}{1+\beta}}^{1} v(1+\beta)\, dv = \frac{3(1+\beta)}{8}.$$

It is comforting to see that:

$$\lim_{(\lambda,\mu)\to(1,1)} p_{2\mathrm{L}} = \hat{p} \quad \text{and} \quad \lim_{(\lambda,\mu)\to(1,1)} \sigma_{2\mathrm{L}} = \hat{\sigma}.$$

We will now demonstrate that the social surplus in Region i, $i \in \{1, 1\mathrm{L}, 2, 3\}$ cannot be greater than $\hat{\sigma} = \frac{3(1+\beta)}{8}$. This proof is simple for Regions 1 and 2. In both these regions, the mass of consumers who buy the product is $\frac{1}{2}$, which is exactly the same mass when $\lambda = \mu = 1$. Furthermore, thanks to the forum, all consumers at $(\lambda, \mu) = (1, 1)$ get full support benefits, that too without incurring any marginal cost. Therefore, the overall consumption and support benefits put together in either Region 1 or Region 2 cannot be any greater than those at $(1, 1)$. Throw in the marginal cost $\chi > 0$ that is incurred in Regions 1 and 2 but not at $(1, 1)$, and it becomes clear that the social surplus in Regions 1 and 2 is dominated by that at $(1, 1)$; formally put, $\sigma_1, \sigma_2 < \hat{\sigma}$.

We will prove $\sigma_{1\mathrm{L}} < \hat{\sigma}$ by contradiction. Suppose that there exists a point (λ, μ) within Region 1L such

that $\sigma_{1\mathrm{L}} \geq \hat{\sigma}$, which is logically equivalent to $\chi \leq \gamma_{1\mathrm{L}}$, where

$$\gamma_{1\mathrm{L}} = \frac{(1+\beta)(1-\mu)\left(3(1-\mu)+\mu(3\beta(1-\mu(1-\lambda))-\lambda)-\sqrt{X_1+X_2}\right)}{3\beta(1+\beta\mu)(1-\mu(1-\lambda))^2},$$
$$X_1 = \frac{9-\mu\left(27+6\lambda-\mu(3+\lambda)(9+\lambda)+\mu^2(3+\lambda)^2-3\beta^2\mu(1-\mu(1-\lambda))^2(3(1-\mu)-\lambda\mu)\right)}{1-\mu}, \text{ and}$$
$$X_2 = \frac{\mu\left(3\beta(1-\mu(1-\lambda))\left(6-3\mu(4+\lambda)+\mu^2(6+\lambda(3-\lambda))\right)\right)}{1-\mu}.$$

We know from Lemma 2 that a valid solution in Region 1L must satisfy $\chi > h_{1\mathrm{L}}$. Now, it can be verified that $h_{1\mathrm{L}} > \gamma_{1\mathrm{L}}$ over the parameter space of interest. In other words, there is no χ that can be small enough to satisfy $\chi \leq \gamma_{1\mathrm{L}}$, yet large enough to satisfy $\chi > h_{1\mathrm{L}}$. This is a contradiction to our earlier supposition that (λ, μ) is a valid solution in Region 1L.

The proof of $\sigma_3 < \hat{\sigma}$ is also by contradiction. Suppose that there exists a point (λ, μ) in Region 3 where, actually, $\sigma_3 \geq \hat{\sigma}$. It can be shown that this condition is logically equivalent to $\chi \leq \gamma_3$, where $\chi = \gamma_3$ is the smaller of the two roots of $\sigma_3 = \hat{\sigma}$. Although a closed form expression for γ_3 exists, it is simply too large to include, even in an appendix. Furthermore, we know from Lemma 2 that a valid solution in Region 3 must satisfy $h_3 \leq \chi \leq g_3$. It can be verified that $h_3 > \gamma_3$ whenever the interval $[h_3, g_3]$ exists, that is, whenever $h_3 < g_3$. Put differently, there can be no χ that satisfies $\chi < \gamma_3$ and $h_3 \leq \chi \leq g_3$ at the same time, which is the desired contradiction.

Finally, it is easy to see that $\sigma_b = \frac{(1+\beta(1-\chi))(3\beta(1+\beta)-\beta\chi(4-\beta))}{8\beta(1+\beta)}$ is a decreasing function of χ. Therefore, its maximum value is obtained when $\chi = 0$; this maximum value is simply $\frac{3(1+\beta)}{8} = \hat{\sigma}$. Therefore, for all $\chi > 0$, σ_b must be less than $\hat{\sigma}$.

We now turn our attention to Region 2L. In this region, the social surplus is always increasing in μ because:

$$\frac{\partial \sigma_{2\mathrm{L}}}{\partial \mu} = \frac{\beta\lambda\left(3+2\beta\chi(1-\lambda)(1+\beta(1-\lambda))+3\beta^2\chi^2(1-\lambda)^2+Y\right)}{8(1+\beta(1-\lambda)(1+\mu(1+\beta)))^2} > 0,$$

where $Y = \beta(1-\lambda)\left(6+3\beta(1-\lambda)+8(1+\beta)(1+\beta(1-\lambda))\mu+4\beta(1+\beta)^2(1-\lambda)\mu^2\right) > 0$. Therefore, the highest social surplus is obtained by setting $\mu = 1$:

$$\hat{\sigma}_{2\mathrm{L}} = \sigma_{2\mathrm{L}}\big|_{\mu=1} = \frac{4\beta^2\lambda(1-\lambda)(1+\beta+2\chi)+3(1+\beta)(1+\beta(1-\lambda)(1-\chi))^2}{8(1+\beta(2+\beta)(1-\lambda))}.$$

It turns out that $\hat{\sigma}_{2\mathrm{L}}$ is concave in λ because

$$\frac{\partial^2 \hat{\sigma}_{2\mathrm{L}}}{\partial \lambda^2} = -\frac{\beta^2(1+\beta)(1+\beta-\chi)(1+\beta+3\chi)}{4(1+\beta(2+\beta)(1-\lambda))^3} < 0.$$

Therefore, solving $\frac{\partial \hat{\sigma}_{2\mathrm{L}}}{\partial \lambda} = 0$, we get $\lambda = \frac{(1+\beta)^2-\sqrt{\frac{(1+\beta)(1+\beta-\chi)(1+\beta+3\chi)}{1+\chi(14-3\chi)+\beta(1+3\chi(2-\chi))}}}{\beta(2+\beta)}$. Of course, for it to be a valid solution, λ must be less than one, which is equivalent to $\chi < \frac{\beta(1+\beta)}{6-2\beta}$. This completes the proof. ■