

DRAWING A LINE IN THE SAND: COMMITMENT PROBLEM IN ENDING SOFTWARE SUPPORT*

Abhijeet Ghoshal

(abhijeet.ghoshal@louisville.edu)

College of Business, University of Louisville, Louisville, KY 40292, USA

Atanu Lahiri

(atanu.lahiri@utdallas.edu)

Jindal School of Management, University of Texas, Dallas, TX 75080, USA

Debabrata Dey

(ddey@uw.edu)

Foster School of Business, University of Washington, Seattle, WA 98195, USA

August 23, 2017

Abstract

We examine the commitment problem faced by a software vendor in ending critical support, in the presence of network security risks. When releasing a new version of a product, in order to drive up its demand, the vendor must cease supporting the old version. However, the vendor's ability to leverage the increased demand can be limited because of a commitment problem. For, when the demand increases and the vendor accordingly sets a higher price, many consumers might opt not to upgrade, creating a situation where stopping security-related support simply becomes too risky. To avoid this risk and any subsequent losses in reputation, the vendor can renege on its earlier decision to stop support. We show that this commitment problem hurts the vendor's profitability and find that the no-commitment equilibrium profit can surprisingly increase with the cost to extend support. Accordingly, we propose a commitment mechanism. Further, the consumer surplus may actually increase if the vendor desists from crossing the proverbial line in the sand and discontinues support as planned.

Keywords: Software support, pricing, security, patching, game theory, commitment problem, analytical modeling, economics of IS.

1 Introduction

Many types of software products, ranging from large enterprise-wide database systems (such as Oracle), operating systems (such as Windows), expensive special-purpose applications (such as Adobe Creative Suite), all the way to smaller desktop utilities (such as Quicken), face a familiar conundrum—unless the vendor abandons supporting older releases, it becomes difficult to encourage

*We thank the review team of *MIS Quarterly* for their valuable feedback.

consumers to upgrade to newer ones, as “complacency from users” invariably kicks in (Kelly 2014). At the same time, however, a software exposes its users to significant security risks when its vendor stops providing critical security patches (Keizer 2014). Given this tension, there is little reason for consumers to take a vendor’s words on the face value anytime it claims that its commitment to ending support is actually firm. It is quite likely that the vendor would say so initially, to encourage consumers to upgrade, only to backtrack later to appease those who hold on to the previous release. Further, the larger the number of consumers who hold on to the previous release, the higher is the likelihood that the vendor is going to backtrack as, then, the potential loss of reputation from a security attack and the accompanying level of outrage are both likely to be of a much greater extent.

A case in point is Microsoft’s proclamation in 2007 that all support for Windows XP would end in April 2014 (Evangelista 2014). This hard deadline did prompt some consumers to upgrade to a higher version (Windows Vista or Windows 7), perhaps because they wanted to avoid security issues that often result from using unsupported software. According to Kelly (2014), “You could hear XP users finally panicking” and upgrading, or even “going the whole hog and buying a new PC.” Not everyone upgraded, though. Despite the looming deadline, as many as one third of all Windows users, which is nearly 30% of all desktop users, continued hanging on to their XP systems (Evangelista 2014), a fraction so sizable that many experts and even governments started pressurizing Microsoft to continue critical security-related support well beyond April 2014 (Reuters 2014, Sun 2014, Timberg and Nakashima 2014). Anticipating a substantial and lasting loss of reputation from security breaches, potential legal liabilities, and a real possibility of severe consumer disenchantment, Microsoft finally buckled. In January 2014, merely three months before the impending deadline, Microsoft effectively backpedaled from its earlier announcement to end all support by declaring that security-related support for Windows XP would, in fact, continue for one more year (Seltzer 2014).

Interestingly, even though Microsoft backtracked on its plan to end support for Windows XP, there are notable examples where the vendor actually remained firm. Adobe, for instance, did not budge from its firm stand that it “will not be fixing the flaws in older editions of Photoshop, Illustrator and Flash Professional—all components of Creative Suite—even though it has rated the bugs as critical” (Keizer 2012). Instead, Adobe insisted that concerned consumers should upgrade

to Creative Suite 6 that addresses the vulnerabilities, despite such tactics causing at that time significant dissatisfaction among its consumers.

Why do different vendors choose different strategies to address the same problem, and what contextual factors influence the eventual choice of strategy? To understand, we must recognize that the dilemma faced by a vendor is often quite complex. Ideally, the vendor wants to pronounce the old software “dead” to attract consumers to the new one—the new version then becomes the only way for consumers to enjoy a secure software. However, a higher willingness on the part of consumers to move to the new version does not necessarily mean more profits for the vendor. For, if the vendor raises the price to exploit this extra willingness, it can run into a *commitment problem*—a higher price would mean fewer consumers upgrading, leading to a situation where the vendor could find it too risky to pursue its earlier decision to end support. In other words, a high price, despite being profitable, is not necessarily an equilibrium outcome unless the vendor can also credibly commit to ending support. The vendor could potentially be caught between the devil of not profiting adequately from the new release—as would be the case if it substantially cuts the price for the upgrade—and a deep sea of consumers who might hang on to the old version if it simply raises that price, suspecting reinstatement of critical support. It is this plight of the vendor that we focus on in this work, in order to better understand the implications of strategic consumer behavior for the vendor’s pricing and support decisions for a software product. In particular, we are interested in addressing the following questions:

- How does the inability to convince consumers of its commitment to end support impact the vendor’s pricing and support strategies? How is the profitability affected?
- Can the vendor devise appropriate mechanisms to credibly signal the firmness of its commitment? When is such a mechanism useful?
- What are the welfare implications of the vendor’s strategy? Could ending support be actually good for consumers?

To answer these questions, we start with a straightforward two-stage model. In the first stage, the vendor names the price of the new version, in addition to announcing whether it intends to abandon support for the old one. In the second, consumers decide whether or not to upgrade.

We identify the conditions under which the optimal solution of this simple two-stage game is not time-consistent, that is, the vendor is left with the option as well as an incentive to renege on its promise to end support after consumers have actually upgraded. Accordingly, we also examine the *no-commitment* game where consumers anticipate the vendor to revisit its earlier decision. This effectively adds a third stage to the game, when the vendor ultimately decides on its support strategy after consumers have made their upgrade decisions. We solve the subgame-perfect equilibrium of this three-stage game. Surprisingly, in the equilibrium, there is a substantial region in the parameter space where the vendor's profit is increasing in the cost involved in supporting the older release. As we explain in detail, this anomalous behavior of the equilibrium profit is not merely surprising but also has important implications for the vendor's overall strategy. In essence, when the support cost is perceived to be low, consumers find the vendor lacking in commitment to end support, which depresses profits. On the other hand, a higher cost simply tells consumers that support is unlikely to be extended, making the commitment stronger and thereby restoring the vendor's profitability.

When the profit is increasing in cost, it is only natural to ask whether the vendor should try to inflate its cost. And, if it so chooses, how would it actually do so? We show that a simple *refund* scheme suffices here. All that the vendor has to do is to issue a contingent refund coupon to every consumer of the new version, which entitles him to a monetary compensation in the event support for the old version is extended. If the vendor ultimately extends support, it will not only incur the cost of patch development but will also be legally obligated to pay a substantial compensation to those who upgraded, raising the effective cost to extend support. This higher cost, in essence, would convince consumers of the vendor's unwillingness to extend support. Accordingly, they would then show a greater willingness to upgrade. In short, the primary benefit of a sufficiently large refund coupon is that it can signal a credible commitment, resulting in a higher revenue.

The issue of consumer and social welfare is also of considerable interest. As mentioned earlier, consumer groups, and sometimes even governments, insist that the vendor should continue supporting older releases in order to prevent large scale security attacks. In contrast, some bloggers and commentators think that crossing the proverbial line in the sand and reinstating support is, in fact a bad idea, as doing so simply causes the penny-pinching users to drag their feet, which, in turn, only prolongs the pain and makes the security environment worse in the long run. As Kelly (2014) bluntly puts it, "you have to be cruel to be kind." A major point of our modeling experiment

is to also put this wisdom to test as well as to establish the conditions for its validity.

2 Literature Review

Several streams of research cross paths with the context we analyze. Of them, particularly important are the ones that concern product upgrades and security patching. We discuss them below. Our research primarily lies in their intersection, in addition to overlapping with the broader economic literature on commitment problem (cf. Fudenberg and Tirole 1991, pp.74–77).

A major focus of the literature on upgrades is on how to price optimally when consumers are strategic and have the option of not buying now in anticipation of a better value later. It is entirely possible that, if consumers expect a product to evolve rapidly in terms of quality, their anticipation can impair a vendor’s ability to practice inter-temporal price discrimination, as an equilibrium becomes impossible to attain (Dhebar 1994). In such a scenario, an equilibrium can be reached only if the vendor precludes upgrade pricing, the strategy of charging a lower price to those who own an older version (Kornish 2001). Of course, the issue of whether upgrade pricing is optimal—or even feasible—is an intriguing question in itself and has also been rigorously examined (Bala and Carr 2009). Likewise, the broader issue of product life-cycle management and upgrade release strategy has also received considerable attention (e.g., Bhattacharya et al. 2003, Mehra et al. 2014, Ramachandran and Krishnan 2008).

Though our emphasis is quite different from the papers above, we share an interesting connection with this upgrade literature. That consumers’ apprehensions about a product’s value can present unique challenges to the vendor is well recognized in this stream. For example, when consumers are apprehensive about the future network size of a product exhibiting positive network effects, somewhat counterintuitively, the vendor may need to introduce a lower quality version first and then offer upgrades (Padmanabhan et al. 1997). Issues concerning forward-compatibility can also create apprehensions about a product’s true value, as consumers of the older version needing compatibility with adopters of the new one are eventually forced to swallow an expensive upgrade (Ellison and Fudenberg 2000). Much like us, Ellison and Fudenberg also identify a commitment problem. In their work, the vendor’s inability to commit to not inflicting compatibility-related pains at a later time leads to depressed profits. Compare this with our vendor, whose inability to commit to ending

support for the older version has similar adverse repercussions on its ability to profit from the new version. In fact, the issue of commitment is often paramount in the context of software upgrades, and it becomes imperative for vendors to assure consumers that they are not making a mistake by deciding to purchase. For software products that are frequently upgraded, a way to accomplish just this and entice consumers to buy the current version is to simply include with it a “New Version Rights” warranty free of charge, that is, to bake in the option to upgrade for free when a new version becomes available (Sankaranarayanan 2007). In a similar vein, we here discuss a mechanism that our vendor can use to overcome its own unique commitment problem.

Moving on to the patching literature, patching has been examined from both the vendor’s standpoint, as well as consumers’ perspectives. The main issue from the vendor’s standpoint is the incentive, or a lack thereof, to offer a secure product. Researchers have identified the fixed-cost nature of patch development and deployment as a plausible incentive to build buggy products (Arora et al. 2006, Choudhary and Zhang 2015). Accordingly, how to incentivize a vendor to either build better products or to provide patches in a timely manner has also been extensively studied (e.g., Arora et al. 2008, 2010, August and Tunca 2011, Kim et al. 2011). Some researchers have further explored the vendor’s incentive to patch illegal users (August and Tunca 2008, Lahiri 2012), while some others have shown that piracy itself can be a motivation to build buggy products (Kannan et al. 2016). We do not consider such motivations or incentives. Rather, we ask whether security support for an older version—which mainly involves providing essential security patches to its users—should be continued and extended beyond the promised time period even after an upgraded software becomes available.

From the consumer’s perspective, the critical issue is when and how often to patch, since frequent patching can be extremely costly (Dey et al. 2015, Ioannidis et al. 2012). There are also related issues, such as whether an option to buy vulnerability information from a marketplace could help consumers better protect themselves, leading to better social outcomes (Kannan and Telang 2005), or whether there could be benefits to aligning consumers’ patch application cycles with the vendor’s patch development/release cycle (Cavusoglu et al. 2008). The issue of incentivizing consumers to patch, in order to prevent them from becoming sources of infection for fellow consumers, has also been examined extensively (August and Tunca 2006).

Finally, we are dealing with a context where consumers without patches adversely impact everyone else by amplifying their security risks, effectively resulting in negative network effects (August and Tunca 2008). There are a number of related papers that study such network effects and/or provide characterizations of various factors that determine security in a user network (e.g., August et al. 2014, Dey et al. 2012, 2014, Galbreth and Shor 2010, Png and Wang 2009, Ransbotham and Mitra 2009). A contribution of this paper is that it brings together security-related negative network effects with commitment problems faced in ending software support.

3 Model Preliminaries

3.1 Context and Overview

In our context, there are two players: a software vendor and a consumer (actually, a set of consumers normalized to a mass of unity). The sequence of actions taken by the vendor is shown in Figure 1. At some point in the past, t_{-1} , the vendor had released a software product, A_{old} ; the perpetual




Figure 1: Timeline for Product Release and Support

license for A_{old} came with a guaranteed support till a future time point, t_1 , which is strictly less than the end of the useful life of A_{old} , denoted t_2 in Figure 1.

At present, that is, at t_0 , the vendor releases a new version, A_{new} , and sets the upgrade price to p .¹ Also, at this time, the vendor may reaffirm or change its support policy for A_{old} . The vendor considers one of the two support strategies: (i) NS , where the vendor, tacitly or explicitly, sticks to its word about ending support at t_1 , and (ii) S , in which the vendor announces an extension of

¹For simplicity, we only consider upgrades in this work, and leave out sales to new consumers, the price for which is often set separately. Most software products sell at a significant discount for the upgrade version. For example, Adobe Creative Suite Version 6 used to sell for \$2,599, but its upgrade price from Version 5 was only \$375 (Keizer 2012). Similarly, the typesetting software PCT_EX Version 6 is priced at \$79, but an upgrade from Version 5 costs just \$39. Such a pricing strategy is widely adopted by vendors of most other software products.

the support window to t_2 .

Consumers then respond to the upgrade offer—some upgrade to A_{new} before t_1 , while the remaining fraction continues with the older version, A_{old} . At t_1 , after consumers have made their upgrade decisions, the vendor can revisit the announced support policy; again, it can either adopt NS and not extend support beyond t_1 , or choose S and extend support to t_2 .

If the vendor forfeits its move at t_1 and sticks to its promise made at t_0 —and if consumers believe that it would—there are no more strategic actions left for either player, and the game is essentially over at t_1 . This situation can be easily abstracted into a discrete-time two-stage game shown in Figure 2, in which the vendor moves first by announcing the price and support strategy. The consumers move second in terms of whether they upgrade or not.




Figure 2: Two-Stage Game on a Discretized Timeline

A word is now in order about why we discretize the timeline. Similar to Sankaranarayanan (2007), the purpose of this work is not to optimize the decision epochs—when to release a product or when exactly to stop support—but to investigate the vendor’s pricing decision and support strategy to identify if it indeed faces a commitment problem when releasing new versions of an old software product. This purpose is amply served by leaving all timing decisions as exogenous to our model. Thus, the time points, t_{-1} , t_0 , t_1 , and t_2 in Figure 1 are all fixed for our analyses. In other words, as shown in Figure 2, the relevant time intervals can be lumped into discrete stages, with model parameters absorbing any aspects that depend on the fixed durations of these intervals. Discretizing the setup in this manner allows us to abstract away unnecessary clutter related to temporal discounting, time-value of use, and other complexities associated with a continuous model and to zoom in onto the commitment issue as our core focus.

3.2 Consumer Model

A consumer's decision to upgrade would depend not only on the price and vendor's support strategy, but also on his own relative valuation for the new version. The consumer may be attracted to the new version for two main reasons:

Functionality Benefits: In addition to offering what is already available in the old one, a newer version may provide, among other things, advanced features and functionalities, a better integration with other applications/environments, a more user-friendly interface, and perhaps most importantly, a longer shelf-life.

Security Losses: If the vendor adopts NS , that is, if it refuses to extend the support beyond t_1 , consumers sticking to the unsupported version, A_{old} , are likely to face increased security breaches and resulting losses. Such losses are not incurred by consumers if they upgrade to A_{new} , or if the vendor adopts S .

We start by considering the functionality benefits. If v_{new} is the consumer's valuation for the new version, and $v_{old} \leq v_{new}$ is the residual value he gets from A_{old} over its remaining useful life, the difference $v = v_{new} - v_{old}$ would represent the functionality-related benefits that he effectively obtains from upgrading. We assume that consumers are heterogeneous and uniformly distributed in their v :²

Assumption 1. *Consumers are indexed by their relative valuation v ; v is distributed uniformly over $[0, 1]$. A consumer knows his v , but the vendor knows only the distribution.*

Let us now consider how consumer v decides on upgrading. When the vendor's adopted support strategy is S , security is not a concern for either version, A_{old} or A_{new} . So, security losses are not considered at all by a consumer, and he upgrades if and only if the following individual rationality (IR) condition is satisfied:

$$v_{new} - p \geq v_{old} \quad \Leftrightarrow \quad v \geq p. \quad (\text{IR-S})$$

²The assumption of uniform distribution for v is simply for analytical tractability. We have verified through extensive numerical analyses that our results hold for most distributional shapes.

In contrast, if the vendor’s strategy is *NS*, the above condition is no longer useful; now, consumers must also weigh security-related concerns resulting from a lack of security support for A_{old} during $[t_1, t_2]$. Put another way, the value of the old version should be even lower to them because of increased security risks. Why does discontinuing support greatly increase such risks? There are two main reasons. First, as new vulnerabilities are discovered from time to time, their remedial patches are not going to be available anymore for an unsupported platform, leaving its users with a heightened risk of getting breached. Second, even if information about a new vulnerability is not in the public domain and is not readily available to them, hackers can always reverse-engineer that information from patches released for the new version, which can then be targeted towards the unsupported older version. Indeed, Microsoft estimated in early 2014 that Windows XP-based machines would be “infected at a rate 66% higher than before patches stopped” (Keizer 2014).

To model this security risk, we follow the literature on negative network effects (August and Tunca 2008). According to this literature, every vulnerable node contributes to security risks faced by other vulnerable nodes in the network, which makes the probability of contracting an infection and getting breached increasing in the number of vulnerable users. Hence, we assume:

Assumption 2. *If the number of consumers using the unsupported version is \bar{v} , the expected security-related loss faced by a consumer without necessary support is $\mu\bar{v}$, where the proportionality constant $\mu > 0$ is essentially the network effect parameter.*

According to Assumption 2, when the support strategy is *NS*, consumer v suffers an additional loss of $\mu\bar{v}$ if he chooses to not upgrade and continues using the old version. Therefore, he now finds upgrading valuable if and only if the following IR condition is satisfied:

$$v_{\text{new}} - p \geq v_{\text{old}} - \mu\bar{v} \quad \Leftrightarrow \quad v + \mu\bar{v} \geq p. \quad (\text{IR-NS})$$

Now, by the very definition of \bar{v} in Assumption 2, consumers with $v \in [0, \bar{v})$ do not upgrade and those with $v \in [\bar{v}, 1]$ do. This clearly means that (IR-NS) must be an equality for the marginal consumer \bar{v} , which immediately leads to $\bar{v} = \frac{p}{1+\mu}$. Combining this with that obtained from a

binding (IR-S), we write:

$$\bar{v} = \begin{cases} p, & \text{for } S, \\ \frac{p}{1+\mu}, & \text{for } NS. \end{cases} \quad (1)$$

The implication of (1) is apparent. Strategy *NS* creates a greater incentive for consumers to upgrade to the new version, decreasing \bar{v} and enhancing the demand for upgrade.

3.3 Vendor’s Problem

Before we can discuss the vendor’s profit-maximization problem, we need to recognize that there is an additional cost associated with strategy *S*. We take cues from prior literature in modeling this support cost as a fixed cost (Arora et al. 2006)—it is essentially the fixed cost incurred in developing additional patches for the extended support window of $[t_1, t_2]$,³ as the marginal cost of distributing them to individual installations is negligible.

Assumption 3. *The vendor incurs an additional fixed cost $c \geq 0$ if it chooses *S*.*

Of course, strategy *NS* avoids this cost, but it has its own risks and associated financial implications for the vendor. One of them is direct. Leaving certain consumers without support could result in them incurring financial losses due to security breaches; the expected value of such losses, aggregated, according to Assumption 2, is $\mu\bar{v}$. The vendor ought to face a proportional financial risk because of legal liabilities directly arising out of security breaches of an unsupported platform. We express this part as $\beta\mu\bar{v}$, where $\beta \geq 0$ is the constant of proportionality.

There could also be an indirect loss to the vendor. That loss has to do with the loss of reputation—a loss of trust, a gradual erosion of brand loyalty—that a vendor might often be concerned about. There is ample anecdotal evidence where vendors of information goods have exhibited, by caving in to consumers’ demands, that they do care about widespread discontent among their consumers.⁴ It is only natural that such reputation loss from an unsupported platform should depend on the size of the unsupported base, \bar{v} . We, therefore, assume that this loss is $\alpha\bar{v}$, where $\alpha \geq 0$ is a constant of proportionality.

³In general the support cost should depend on $(t_2 - t_1)$, the size of the additional support window. Since the time points are fixed in our analysis, the size of the window is also fixed and so is the support cost.

⁴As a few indicative examples, consider the following: Sony took away DRM restrictions on its music sales, Microsoft brought back the start button in Windows 8.1, and Apple apologized to its consumer base and allowed Google Maps on its devices after the debacle of melting bridges and roadways in Apple Maps.

What is interesting about this reputation loss or consumer backlash, $\alpha\bar{v}$, is that it figures only in the vendor's calculations, perhaps as lost future revenue, but does not enter into the consumer's utility function, because it is really not a cost the consumer bears in practice. The "bad press" or the "bad taste in mouth," the resulting erosion of trust or brand loyalty, and disillusionment about a product are all a part of consumers' overall learning process about a product or a brand, and may figure, implicit- or explicitly, in his future purchase decisions, but is certainly not pertinent to him at present. Taken together, we can write:

Assumption 4. *The total financial loss faced by the vendor, when it chooses NS, is $(\alpha + \beta\mu)\bar{v}$, where $\alpha \geq 0$ and $\beta \geq 0$.*

The motivation behind this abstraction is straightforward. What we want to show is that, as long as there are security-related network effects, that is, as long as $\mu > 0$, our results hold for a general formulation of the vendor's loss. In this formulation, either of the two parameters, α and β , can be zero or arbitrarily large. Thus, our setup remains valid irrespective of the extent to which the effects of network security risks spill over from consumers to the vendor. When $\alpha = 0$ or $\beta \gg \alpha$, the vendor's loss is primarily from the spill-over. In contrast, when $\beta = 0$ or $\alpha \gg \beta$, the contribution of the spill-over towards the vendor's loss is negligible. Our analyses apply to both these extremes as well as all points in-between.

For ease of exposition, we now define the vendor's *loss rate* as the normalized financial loss borne by it:

$$\gamma = \frac{\alpha + \beta\mu}{1 + \mu}.$$

In this paper, we focus on the case of $\gamma < 1$; the case of $\gamma \geq 1$ is uninteresting as, then, the vendor no longer considers NS because of the prohibitive nature of the expected financial losses.

Recall that the demand for upgrade is $(1 - \bar{v})$. Using \bar{v} from (1), we can write the vendor's profit, net of appropriate costs and losses, as:

$$\pi = \begin{cases} p(1 - p) - c, & \text{for } S, \\ p\left(1 - \frac{p}{1+\mu}\right) - \gamma p, & \text{for } NS. \end{cases} \quad (2)$$

This profit function is concave in p in each case. Using appropriate first order conditions, we can

obtain the optimal price and the corresponding profit in each case:

$$p^* = \begin{cases} p_S = \frac{1}{2}, & \text{for } S, \\ p_{NS} = \frac{(1-\gamma)(1+\mu)}{2}, & \text{for } NS, \end{cases} \quad \text{and} \quad \pi^* = \begin{cases} \pi_S = \left(\frac{1}{4} - c\right), & \text{for } S, \\ \pi_{NS} = \frac{(1-\gamma)^2(1+\mu)}{4}, & \text{for } NS. \end{cases} \quad (3)$$

Clearly, when c is large enough to make $\pi_{NS} > \pi_S$, it is optimal for the vendor to choose NS . Likewise, the converse is true when c is not large. Solving $\pi_{NS} = \pi_S$, we get $c = \frac{1-(1-\gamma)^2(1+\mu)}{4} = c_1$, which leads to:

$$\text{Optimal Strategy} = \begin{cases} \text{Choose } S \text{ and } p = p_S, & \text{when } c \leq c_1, \\ \text{Choose } NS \text{ and } p = p_{NS}, & \text{otherwise.} \end{cases} \quad (4)$$

Though optimal, the solution in (4) is not necessarily time-consistent since, at t_1 , it leaves the vendor with the option of revisiting the issue of support and, under certain conditions, with an incentive to extend support to t_2 . Formally,

Proposition 1. *Let $c_1 = \frac{1-(1-\gamma)^2(1+\mu)}{4}$ be as above, and let $c_2 = \frac{\gamma(1-\gamma)(1+\mu)}{2}$. If $c_1 < c < c_2$, the optimal solution in (4) is not time-consistent—specifically, the vendor has an incentive to revert to S after consumers have made their upgrade decisions.*

Proposition 1 essentially means that consumers cannot blindly trust the vendor to follow through on its earlier announcement about discontinuing support. This is because once consumers have spent on the new version, the vendor can simply backtrack and extend support, thereby mitigating its financial losses. This possibility that the vendor might renege on its earlier promise can have serious consequences. In fact, the vendor could face a commitment problem, and its inability to commit to NS could spell adverse impacts on its profits. We discuss this issue and the ensuing *no-commitment equilibrium* in detail in our next section.

This commitment problem, as well as the issue of time-inconsistency, arises in many practical situations. For example, macro-economists have shown that a government—even when it should be interested in holding inflation to zero—can be left with an unexpected incentive to raise the inflation level after private citizens have made their output decisions (Mankiw 1988, pp. 441–443). Thus, a zero inflation, though optimal to begin with, is not time-consistent and, therefore, also not “an equilibrium of the game without commitment” (Fudenberg and Tirole 1991, p.76).

Software vendors, too, run into similar commitment problems when deciding whether to launch newer versions of their products (Sankaranarayanan 2007); as shown there, even when it is optimal up front for a vendor to not consider newer versions of a software, the vendor might have an unexpected incentive to do exactly the opposite at a later point in time. Further, if consumers can suspect a reversal in the vendor’s strategy and start expecting a new version at a later date, “no-new-versions” ceases to be an equilibrium in the absence of commitment.

4 No-Commitment Equilibrium

When consumers do not trust the vendor and expect it to choose its *actual* support strategy only after their upgrade decisions, their decision at the second stage could be quite different. The resulting strategic interactions can be captured in a discrete-time three-stage game where the vendor makes an additional third move at t_1 . Consequently, the timeline in Figure 1 reduces to Figure 3. A point to note in Figure 3 is that it only considers the case in which the vendor announces NS




Figure 3: Three-Stage Game on a Discretized Timeline

at t_0 . This is because there is no commitment issue if the vendor announces S ; accordingly, we exclude that possibility from further considerations.⁵

Comparing Figures 2 and 3, it is apparent that the difference between the two cases is that the vendor’s incentive to revisit the support strategy essentially adds a third stage to the game, rendering the announcement of NS at stage 1 irrelevant and changing the timing of the actual support decision. Earlier, this decision happened in stage 1; now, it happens in stage 3. This realignment of the timeline is akin to what happens in the case of sequential duopoly games involving

⁵Theoretically, the vendor’s strategy space—its announcement at t_0 versus its action at t_1 —consists of four pairs: $\{\langle S, S \rangle, \langle S, NS \rangle, \langle NS, S \rangle, \langle NS, NS \rangle\}$, with the second and third choices reflecting a potential commitment problem. However, in reality, the vendor can never promise support (S) and then renege (NS), as that would be grossly illegal—in fact, bordering on consumer fraud—and would expose the vendor to huge legal liabilities and irreparable reputation loss. In addition, the vendor gains nothing by announcing S at t_0 ; it only loses some revenue because some consumers will not upgrade now. Thus, we are practically left with only two choices: $\{\langle NS, S \rangle, \langle NS, NS \rangle\}$.

quantity competition (Osborne 2004, p.192): Absent a commitment that the leader will not readjust its production quantity after the follower has made its decision, the follower starts acting as if it is the real first mover, causing the leader to become the second mover for all practical purposes. The unique aspect in our case is that the vendor—the leader—makes two decisions. One of these two, the pricing decision, still happens at stage 1, but the actual support decision now moves to a later time. Thus, our vendor too loses its first-mover status, albeit only partially.

As is customary, we now traverse the timeline in Figure 3 backward, in order to obtain the desired equilibrium. Having already pocketed all its revenue at stage 2, now, at stage 3, the vendor incurs only one of the following two costs: (i) from Assumption 3, a support cost c when choosing S , or (ii) a financial loss that results from NS , $(\alpha + \beta\mu)\bar{v} = \gamma\bar{v}(1 + \mu)$ according to Assumption 4. So, at stage 3, the vendor has an incentive to stick to NS only when:

$$\gamma\bar{v}(1 + \mu) \leq c \quad \Rightarrow \quad \bar{v} \leq \frac{c}{\gamma(1 + \mu)}. \quad (5)$$

Moving back to stage 2 of the game, we consider the upgrade decision facing our consumers. Consumers should expect the vendor to stick to NS if \bar{v} resulting from such a strategy indeed abides by (5). Recall from (1) that $\bar{v} = \frac{p}{1+\mu}$ for NS . Substituting this into (5) leads to:^{6,7}

$$p \leq \frac{c}{\gamma}. \quad (6)$$

Condition (6) implies that consumers would have reasons to suspect that the vendor, notwithstanding its earlier announcement, will revert to S if the price exceeds the critical fraction $\frac{c}{\gamma}$.

Finally, moving back to the first stage of the game, we look at the vendor's pricing decision. Our vendor knows that consumers would anticipate NS when (6) is satisfied, and S otherwise. So,

⁶Here, instead of using $\bar{v} = \frac{p}{1+\mu}$ for NS , could we have used $\bar{v} = p$ for S , resulting in a threshold of $\frac{c}{\gamma(1+\mu)}$? Of course, if $p \leq \frac{c}{\gamma(1+\mu)}$ or if $p > \frac{c}{\gamma}$, there is no ambiguity with respect to consumers' anticipation. So, the question becomes one of how consumers would behave when $\frac{c}{\gamma(1+\mu)} < p \leq \frac{c}{\gamma}$. If a consumer continues to anticipate NS based on $p \leq \frac{c}{\gamma}$, there is nothing more to do. On the other hand, if the consumer anticipates S based on $p > \frac{c}{\gamma(1+\mu)}$ and decides to not upgrade, the vendor can still stick to NS at stage 3, thereby making this consumer face imminent security threats and forcing him to upgrade in a rush. Since consumers know that the vendor can do this, and since the vendor knows that consumers know that the vendor can squeeze consumers this way, the only possible outcome is that all consumers would only anticipate NS in the region $\frac{c}{\gamma(1+\mu)} < p \leq \frac{c}{\gamma}$, which means that the threshold of $\frac{c}{\gamma(1+\mu)}$ is practically irrelevant.

⁷We assume that all the cost associated with NS is incurred at t_1 . It is, however, possible that a portion of this cost may be incurred even earlier and may be sunk at t_1 . If that happens, the denominator in (6) will decrease accordingly, making (6) easier to satisfy and giving the vendor more incentive to keep its commitment.

the vendor solves:

$$\pi = \begin{cases} p(1-p) - c, & \text{if } p > \frac{c}{\gamma} \text{ (S)}, \\ p\left(1 - \frac{p}{1+\mu}\right) - \gamma p, & \text{if } p \leq \frac{c}{\gamma} \text{ (NS)}. \end{cases} \quad (7)$$

The main difference between (2) and (7) is that, in (2), the support and no-support strategies were not constrained by the threshold $\frac{c}{\gamma}$. Maximizing the profit in (7) with respect to p leads to the following partitioning of the parameter space:

Proposition 2. Let $c_1 = \frac{1-(1-\gamma)^2(1+\mu)}{4}$ and $c_2 = \frac{\gamma(1-\gamma)(1+\mu)}{2}$ be as before. Further, let us define $c_3 = \frac{\gamma(1+\mu)}{2} \left(1 - \sqrt{\frac{\mu}{1+\mu}}\right)$. The equilibrium of the no-commitment game can then be characterized as follows.

- **Support Region (S):** If $c \leq c_3 < c_2$, or if $c_2 \leq c_3$ and $c \leq c_1$, the vendor continues support.
- **No-Support Region (NS):** If $c_3 < c_2 \leq c$, or if $c_2 \leq c_3$ and $c > c_1$, the vendor discontinues support.
- **Limit Region (NS-L):** If $c_3 < c < c_2$, the vendor uses a “limit price”—mathematically a corner solution—while discontinuing support.

Proposition 2 is better visualized in Figure 4 which illustrates, for $\alpha = 1$ and $\beta = \frac{1}{4}$, how the entire (γ, c) space partitions into three regions, with each region representing the set of parameter values corresponding to a particular outcome. In the first outcome, denoted S for support, the optimal profit is simply the maximum of $\pi_S = p(1-p) - c$, obtained from the appropriate first order condition. Likewise, in the second, denoted NS for no support, it is the interior maximum of $\pi_{NS} = p\left(1 - \frac{p}{1+\mu}\right) - \gamma p$. In other words, the optimal prices in these two regions are p_S and p_{NS} , respectively, as given by (3). Interestingly, there is also a third region, a corner solution in which the vendor simply sets the price at a limiting value of $p_L = \frac{c}{\gamma}$, instead of choosing either interior solution; in the figure, we refer to this last possibility as $NS-L$, with the letter L signifying *limit*.

Figure 4 also makes it easy to compare Proposition 2 with Proposition 1. According to Proposition 1, the vendor faces a commitment problem for $c_1 < c < c_2$; this is essentially the region below the c_2 curve in the figure, except for the portion that is also below the c_1 curve. As Proposition 2 shows, the impact of this commitment problem is actually two-fold. Between c_3 and the dashed

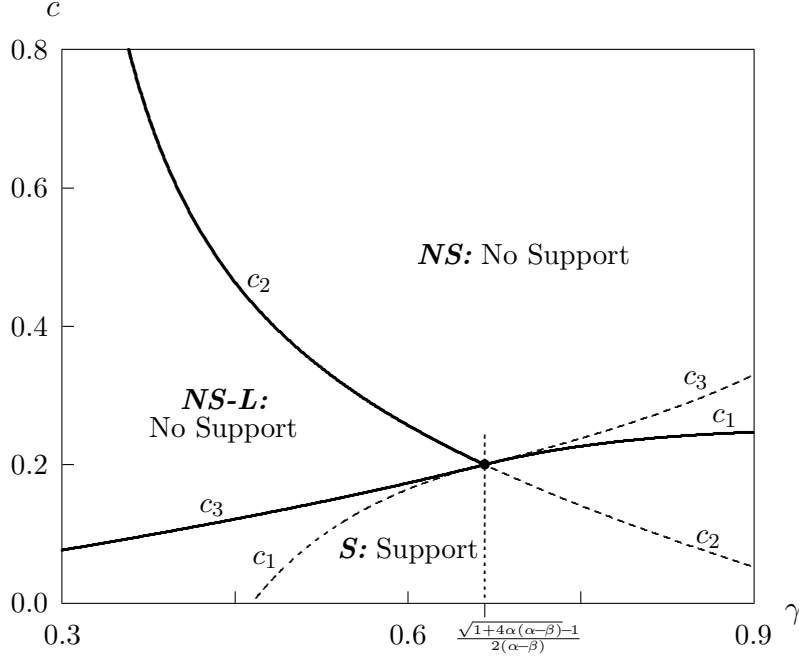


Figure 4: Relevant Partitions of the (γ, c) Space; $\alpha=1$, $\beta=\frac{1}{4}$

portion of c_1 , unable to commit to NS , the vendor simply gives up on NS as if that is not even an option. On the other hand, between c_3 and c_2 , the vendor forces no-support by means of a corner solution, that is, by choosing $p = \frac{c}{\gamma}$ to barely satisfy the condition for NS in (7).

In Figure 5, we plot the equilibrium price, which is given by:

$$p^* = \begin{cases} p_S = \frac{1}{2}, & \text{in Region } S, \\ p_{NS} = \frac{(1-\gamma)(1+\mu)}{2}, & \text{in Region } NS, \\ p_L = \frac{c}{\gamma}, & \text{in Region } NS-L, \end{cases} \quad (8)$$

as a function of γ and c , again for $\alpha = 1$ and $\beta = \frac{1}{4}$. There are several important observations that can be made from this figure. First, when the support cost, c , is low, the vendor decides to be in the support region, Region S , and it simply charges the traditional monopoly price of $p_S = \frac{1}{2}$. In this region, the vendor may as well admit ahead of time that it would eventually extend support to all consumers. For, even if the vendor keeps repeating the tale that it is firmly committed to ending all support, consumers would duly call its bluff, ultimately shifting the equilibrium to Region S . Thus, everywhere in Region S —including the portion above the dashed part of c_1 , where supporting is




Figure 5: Equilibrium Price as a Function of γ and c ; $\alpha=1$, $\beta=\frac{1}{4}$

not optimal according to (4)—the profit-maximizing vendor would extend support regardless of its earlier announcements. The sheer absence of credible commitment does make the possibility of a vendor going back on its word and reverting to S very real.

Second, there is a significant portion of the parameter space (NS and $NS-L$), where the higher cost of support makes it imperative for the vendor to end support for the old version. Especially interesting is the $NS-L$ region. Though it may seem surprising at the first glance, the $NS-L$ strategy is both practically relevant and meaningful. Evidently, when $c_1 < c < c_2$, the vendor does not have enough leeway to pursue the NS strategy, despite it being the more profitable one. This is because a high p_{NS} works like a trap for the vendor, driving many consumers away from the upgrade, raising in turn doubts about its own commitment to truly end critical support. As explained above, in the region between c_3 and the dashed part of c_1 , the vendor's response is to simply revert to S . In contrast, in the region between c_3 and c_2 , the vendor finds a way to avoid falling into the trap. In fact, it does so by holding the price at $p_L = \frac{c}{\gamma}$, a rather low level that is clearly below p_{NS} ; see

Figure 5. The implication is immediate. Between c_3 and c_2 , a vendor can, and should, use its price to signal a commitment to *NS*. If consumers are convinced that a sufficient fraction of them would be eventually upgrading, they would indeed act as if the vendor would forsake support for good. Viewed another way, a vendor who fails to recognize this connection between the upgrade price and the support strategy could end up choosing a price that is suboptimally high.

There is, in fact, a real possibility that Microsoft walked into this trap when announcing the death of Windows XP. Recall from our earlier discussions that nearly a third of all Windows users continued using XP even as the April 2014 deadline loomed large (Evangelista 2014), and the prevalent wisdom at that time was also that Microsoft will eventually bend and “Windows XP isn’t going anywhere” (Keizer 2014). A quick examination of Windows’ prices does lend some credibility to our assessment that a high upgrade price might have led to this situation. For example, when Microsoft announced the release of Windows Vista, the retail price for an upgrade to the Business version was \$199; this price was quite high, especially when compared with the full retail price of \$299 that a first-time Windows buyer was expected to pay. In other words, the price for the upgrade was about 67% of the full price. Microsoft stuck to the very same strategy when pricing Windows 7 in 2009. At that time, the XP to 7 upgrade for the Ultimate (respectively, Professional) version was priced at \$219 (respectively, \$199), which is 69% (respectively, 67%) of the full price of \$319 (respectively, \$299). Compare this with Adobe’s support and pricing strategy. Adobe’s price for the Creative Suites 6 (CS6) upgrade was only \$375 (Keizer 2012), a mere 14% of the full price (\$2,599). And, in line with our findings, when the moment to choose arrived, unlike Microsoft, Adobe saw little need to go back on its earlier decision to end support (Keizer 2012). In other words, our analysis suggests that, if a mistake were at all made by Microsoft, it was certainly not made in 2014, as by that time Microsoft’s hands were practically tied by an earlier mistake which, in all likelihood, was first made in 2007 at the time of the Vista release and then repeated in 2009 when Windows 7 was announced.

Is there a third possibility? Is it at all possible that the support cost, c , for XP was truly high, but the consumers grossly underestimated it? Our observations of the events leading up to Microsoft’s decision to extend support indicate that, as the April 2014 deadline approached, Microsoft became quite desperate in its attempts to label XP as really old and not fit for the technically advanced world of today. Experts siding with Microsoft even went to the extent of

characterizing the continued use of XP as “taking a World War II-era tank and putting it in the battlefields of Iraq or Afghanistan today” (Evangelista 2014). Therefore, it is quite possible that there existed severe underestimation on the part of consumers and the rightful equilibrium in the NS region could not be achieved. Although we do not explicitly model such systematic underestimation, the results and discussions in the next section do provide a way to address this issue through a commitment mechanism.

5 Profit and Commitment Mechanism

How is the vendor’s profit affected by the commitment problem? To investigate this, we substitute the price from (8) into (7) and obtain the profit in the no-commitment equilibrium:

$$\pi^* = \begin{cases} \pi_S = \left(\frac{1}{4} - c\right), & \text{in Region } S, \\ \pi_{NS} = \frac{(1-\gamma)^2(1+\mu)}{4}, & \text{in Region } NS, \\ \pi_L = c \left(\frac{1}{\gamma} \left(1 - \frac{c}{\gamma(1+\mu)}\right) - 1\right), & \text{in Region } NS-L. \end{cases} \quad (9)$$

Proposition 3. *If $c_1 < c < c_2$, the vendor’s profit is adversely impacted because of the commitment problem. Further, the profit in the no-commitment equilibrium is non-monotonic in c ; it is decreasing in c in Region S , increasing in Region $NS-L$, and flat in Region NS .*

Recall that the vendor prefers NS when $c > c_1$, but also faces a commitment problem when $c_1 < c < c_2$; see Proposition 1. Proposition 3 essentially means that, when the vendor is unable to commit to NS , it simply ends up with a profit of π_S or π_L , both of which are less than π_{NS} . It is thus apparent that the commitment problem leaves the vendor worse off. The lesson could not be any clearer. The ability to readjust the support strategy after consumers’ decisions may have a natural appeal to some managers, but it is this very ability that eventually boomerangs and causes the vendor to lose money. Though a bit surprising, such is typically the irony of the fate of an economic agent facing a commitment problem—when the optimal solution turns out to be inconsistent, the consistent ones invariably turn out to be suboptimal (Kydland and Prescott 1977, pp. 474–475).

Clearly, addressing the commitment problem is paramount for the vendor. To gauge how the

vendor might be able to do so, we now plot the no-commitment profit as a function of γ and c , again for $\alpha = 1$ and $\beta = \frac{1}{4}$. Interestingly, when γ is low, the no-commitment equilibrium profit is




Figure 6: Equilibrium Profit as a Function of γ and c ; $\alpha=1$, $\beta=\frac{1}{4}$

not necessarily decreasing in the cost of continuing support, c . As Proposition 3 points out, only when the strategy in use is S , it declines as expected. However, in Region $NS-L$, it is surprisingly increasing in c . Finally, it becomes flat upon crossing over to Region NS , where the cost to support has no bearing on the equilibrium profit. For relatively higher values of γ , however, Region $NS-L$ is non-existent; as a result, the profit is initially decreasing in Region S and then flat afterwards in Region NS . A key lesson here is that, when the perceived damage from security attacks is high but the vendor faces a relatively smaller indirect impact, the vendor could surprisingly prefer its patch development costs to be high. Specifically, it would prefer a c at least as large as c_2 to one in (c_1, c_2) .

That a vendor can at all prefer higher costs under some circumstances is not only fascinating but is also of significant relevance. Such anomalous behavior of profit is, in effect, an indication

that the vendor can overcome the adverse effect of the commitment problem by suitably inflating its cost to extend support. The obvious question here is: how can the vendor credibly “inflate” the cost?⁸ Among many possibilities, the most intuitive one ought to be to include a contingent refund coupon—some sort of a money-back guaranty—into the sales contract for A_{new} . In the event support is extended for A_{old} , this coupon would automatically entitle those who have upgraded to a monetary compensation. In fact, software makers already use such refunds for other purposes (e.g., Protalinski 2014), and existing literature also recommends similar mechanisms that can be easily incorporated into sales contracts (Sankaranarayanan 2007).

To design such a mechanism, note that, in the event support is extended, the vendor is obligated to pay out the promised refund, denoted r , to each consumer who produces the coupon. In other words, the effective support cost borne by the vendor now becomes $(c + r(1 - \bar{v}))$. In order to address the adverse impact of the commitment issue, all the vendor has to do is to choose an r so that $(c + r(1 - \bar{v}))$ becomes larger than c_2 . Essentially, the purpose is to burn the proverbial bridge (Fudenberg and Tirole 1991, p.75), simply to convince consumers that there is no way for the vendor to retreat once it has promised to end support. The following result discusses how large the refund needs to be for it to be suitable for this specific purpose.

Proposition 4. *If $c_1 < c < c_2$, a refund coupon of $r \geq \rho = \frac{\gamma(1-\gamma)(1+\mu)-2c}{1+\gamma}$, payable to each buyer of the new version in the event support is extended for the old version, is sufficient to fully address the commitment problem.*

Proposition 4 simply tells us that the size of the refund must be large enough for it to serve as a commitment mechanism. Interestingly though, as long as it is larger than ρ , its actual size has no eventual impact on the optimal profit, which remains π_{NS} . This is because the vendor does not owe any refund as long as it sticks to the committed strategy of NS . Evidently, what matters here is the anticipated cost of choosing S , and not the realized cost in optimality. If consumers anticipate S to be sufficiently costly from the vendor’s perspective, they would stop betting on an extended support and instead show a greater willingness to pay for the upgrade, restoring fully the vendor’s pricing power in the presence of negative network effects.

⁸We write the word “inflate” within quotes, because the support cost is simply the cost to provide patches for the extended support window and cannot really be inflated in that sense. Essentially, the vendor makes strategy S more costly by imposing some kind of enforceable penalty on itself. The key here is that not only must the penalty be enforceable if the vendor chooses S , but consumers must also believe it is indeed so.

Curiously, the above coupon-based scheme can be easily implemented even when there is sufficient uncertainty about the value of c , that is, when information asymmetry between the vendor and consumers makes it difficult for the former to convey the true cost of patch development to the latter. Facing such a situation, the vendor can simply presume the worst-case scenario of $c = 0$ and set the refund coupon at $\frac{\gamma(1-\gamma)(1+\mu)}{1+\gamma}$, to fully assure consumers that the commitment to ending support is indeed firm. It is only by signaling how significant the cost to extend support could be that our vendor can convince the consumer base about the firmness of its commitment.

6 Welfare

We now turn our attention from private profit to public welfare. The consumer surplus generated from the upgrade, net of the losses suffered from security attacks, can be defined as follows:

$$CS = \begin{cases} \underbrace{\int_{\bar{v}}^1 (v-p)dv}_{\text{Gain from upgrade}}, & \text{for } S, \\ \underbrace{\int_{\bar{v}}^1 (v-p)dv}_{\text{Gain from upgrade}} + \underbrace{\int_0^{\bar{v}} (-\mu\bar{v})dv}_{\text{Loss from security attacks}}, & \text{for } NS \text{ and } NS-L. \end{cases} \quad (10)$$

As before, \bar{v} depends on the vendor's strategy according to (1), and $\mu\bar{v}$ captures the effect of security-related breaches as felt by each user of the old version. As appropriately reflected in (10), this adverse impact, totaling $\int_0^{\bar{v}} (-\mu\bar{v})dv$, is not a consideration when the consumers sticking to the old version are provided the necessary support.

To compute the surplus value at equilibrium, we now need to substitute p^* from (8) into (10). This leads to:

$$CS = \begin{cases} \phi_S = \frac{1}{8} & \text{in Region } S, \\ \phi_{NS} = \frac{(1+\gamma)^2 - 4\mu(1-\gamma)}{8}, & \text{in Region } NS, \\ \phi_L = \frac{c^2 - \gamma(1+\mu)^2(2c-\gamma)}{2\gamma^2(1+\mu)^2}, & \text{in Region } NS-L. \end{cases} \quad (11)$$

A closer examination of this consumer surplus reveals a curious result. It is clear that, when the vendor adopts the support strategy, the consumer surplus is $\frac{1}{8}$, exactly what one would expect in a traditional monopoly. Now, consider a situation where the vendor decides—perhaps somewhat irrationally—to be always “kind,” and extends support even when it is not optimal to do so. It is natural to expect that, in such a case, the consumers would all be happier—getting free support ought to reduce their security risks and make them all better off. Such an intuition, however, is not necessarily correct; there is a significant part of the parameter space where the surpluses from *NS* and *NS-L* turn out to be higher than $\frac{1}{8}$. Evidently, under certain circumstances, even if the vendor becomes “cruel” and behaves as a profit-maximizer, it may actually end up providing a better surplus to its consumers. Thus, despite the grave concerns expressed by many consumer advocates and even some governments, there may indeed be some merit to the claim that, at times, “you have to be cruel to be kind” (Kelly 2014):

Proposition 5. *At a relatively large loss rate, the consumer surplus from *NS* or *NS-L* dominates that from *S*. Specifically, $\phi_L > \phi_S$ if $\gamma > \theta(c, \mu)$, and $\phi_{NS} > \phi_S$ if $\gamma > \Theta(\mu)$, where $\theta(c, \mu) = \frac{2c(2(1+\mu)+\sqrt{1+4\mu(2+\mu)})}{3(1+\mu)}$ and $\Theta(\mu) = \sqrt{1+4\mu(2+\mu)} - (1+2\mu)$.*

The implication of Proposition 5 is as follows. When γ is above the threshold of $\Theta(\mu)$, in the no-commitment equilibrium, the vendor’s interest is surprisingly aligned with consumers’ interests in both Regions *NS-L* and *NS*. Further, if the vendor does use a commitment mechanism such as a refund coupon to implement *NS*, that would also be in line with consumers’ interests. On the other hand, when γ is below $\Theta(\mu)$, consumers would no longer prefer *NS* to *S*, but it is still possible that they would prefer *NS-L* to *S* as long as $\gamma > \theta(c, \mu)$.

The intuition as to why consumers surprisingly benefit from *NS* in certain situations is also quite interesting. Although it is true that those sticking to the older version are harmed when support is discontinued, the benefits obtained from more consumers switching to a better product does sometimes outweigh that negative impact. This is precisely what aligns the collective interest of consumers with that of the vendor. Specifically, if γ is somewhat large compared to μ —a smaller μ also implies a smaller $\Theta(\mu)$ —in regions where the vendor finds ending support to be preferable, the consumers, taken as a whole, would also find so. Essentially, a larger γ means a lower p_{NS} , which, in turn, means more consumers upgrading as well as more surplus for each of them upgrading. At

the same time, a lower μ means a smaller security-related impact on those not upgrading, once again making *NS* more appealing to consumers. Therefore, governments and consumer advocates who clamor against discontinuation of support by pointing to security concerns should consider moderating their views on this subject. In fact, some moderation would be particularly desirable in situations where the negative network effect, as represented by μ , is not severe to begin with.

We end this section with a couple of observations on social welfare, which is the sum of the vendor's profit and the consumer surplus. First, under the conditions mentioned in Proposition 5, social welfare also benefits from *NS*. This is expected because, then, both the vendor and consumers gain from *NS*. Actually, their aggregate welfare is higher in a wider region than the conditions in Proposition 5 might suggest, as the aggregate can be higher even when the consumer surplus is not so. This is another reminder that the much maligned decisions of the software vendors can, in fact, be good for the society. Second, the anomalous behavior of profit with respect to c , which is illustrated in Figure 6, carries over to social welfare as well, at least to an extent, making the collective welfare surprisingly higher in situations where the patches are, in fact, costlier to develop. Although somewhat surprising that social welfare can increase with costs incurred by any economic actor, this result is in line with prior literature, where as well an increase in a cost term faced by one actor leads to a strategic response by another, which turns out to be favorable to collective welfare. For example, Mehra et al. (2012) find that, when consumers face a higher switching cost, the software vendor responds with a lower first-period price, leading to unpredictable gains in the overall surplus. In our case, too, when the vendor faces a higher support cost, more consumers respond positively to the upgrade offer, and surprising gains ensue.

7 Robustness Checks

In this section, we look at a few model variations in order to verify robustness of our findings.

7.1 Imperfect Consumer Information

We first revisit the issue of consumers' perfect knowledge of the critical fraction, $\frac{c}{\gamma}$. Recall that, in the no-commitment game in Section 4, in order for consumers to correctly anticipate *S/NS*, they had to compare whether the vendor-announced price, p , exceeds this critical fraction; see (6).

If the vendor cannot rely on consumers to behave in that fashion, its pricing strategy at stage 1 would be different from the equilibrium we found in Section 4. Naturally, questions arise regarding how critically our results depend on this assumption. For, in practice, it is hard to imagine that consumers would have perfect information about the vendor's support cost or its loss rate. In this section, we examine what would happen if, in fact, they do not.

In economic models, a lack of perfect information is typically captured by a probability distribution around the true value of the quantity of interest, which is $\frac{c}{\gamma}$ in our context. Such a distribution simply means that consumers are heterogeneous in h , their assessment of $\frac{c}{\gamma}$, but the mean of their assessments is $\bar{h} = \frac{c}{\gamma}$, same as the actual value faced by the vendor. Any distribution symmetric about $\frac{c}{\gamma}$ will do; we consider a uniform prior:

Assumption 5. *Consumers are heterogeneous in their assessment, h , which is uniformly distributed over $\left[0, \frac{2c}{\gamma}\right]$. Consumers do not know this distribution, but the vendor does.*

When consumers are heterogeneous in this fashion, irrespective of the true value of $\frac{c}{\gamma}$, one group of consumers would, correctly or incorrectly, find that their h satisfies (6) after p is announced, that is, $p \leq h$; at the same time, again correctly or incorrectly, the remaining consumers would not satisfy (6), that is, $p > h$. Of course, consumers do not know the distribution of h , and they do not know whether they are correct or not. Thus, in this extended setup, they face the possibility of making a mistake if their anticipation of NS (or S) turns out to be not true, that is, the upgrade turns out to be less (or more) attractive than originally anticipated. Two cases are possible:

Case 1 ($p > \frac{c}{\gamma}$): It can be shown that the vendor would ultimately extend support in this case, and all consumers with $h \in [0, p)$ would correctly infer so; because $h \sim U\left[0, \frac{2c}{\gamma}\right]$, there will be $\eta = \frac{p\gamma}{2c}$ such consumers. However, there will also be a fraction $(1 - \eta)$ of consumers who would mistakenly believe NS . Of them, a $\left(p - \frac{p}{1+\mu}\right) = \frac{\mu p}{1+\mu}$ fraction—those with $v \in \left(\frac{p}{1+\mu}, p\right)$ —may feel cheated, since they would not have upgraded if they had perfect information.

Case 2 ($p \leq \frac{c}{\gamma}$): In this case, the vendor's strategy would be NS . Although all consumers with $h \geq p$ would correctly guess that strategy, a fraction $\eta = \frac{p\gamma}{2c}$ of consumers would mistakenly assume S . Further, a $\left(p - \frac{p}{1+\mu}\right) = \frac{\mu p}{1+\mu}$ fraction of them will not upgrade because of their wrong assessment. When the vendor sticks with NS at stage 3, however, these consumers will face imminent security threats and will be forced to upgrade in a haste.

We are now ready to revise (7) as follows.

$$\pi = \begin{cases} p(1-p) - c, & \text{if } p \geq \frac{2c}{\gamma} \text{ (} S \text{),} \\ \eta p(1-p) + (1-\eta)p\left(1 - \frac{p}{1+\mu}\right) - c, & \text{if } \frac{c}{\gamma} < p < \frac{2c}{\gamma} \text{ and } p < 1 \text{ (} S' \text{),} \\ (1-\eta)p\left(1 - \frac{p}{1+\mu}\right) - c, & \text{if } \frac{c}{\gamma} < p < \frac{2c}{\gamma} \text{ and } p \geq 1 \text{ (} S'' \text{),} \\ p\left(1 - \frac{p}{1+\mu}\right) - \gamma p, & \text{if } p \leq \frac{c}{\gamma} \text{ (} NS \text{),} \end{cases} \quad (7')$$

where, as defined above, $\eta = \frac{p\gamma}{2c}$. Even though there are four possible cases in (7'), it can be shown after some algebra that case S'' can never happen in optimality and can be dropped from further considerations.⁹ Maximizing this new profit, we can obtain the equilibrium outcome. The cumbersome details themselves are of little value to the core issue in this paper and are omitted in favor of a visual illustration in Figure 7, which shows the equilibrium profit for the same parameter values as in Figure 6, that is, $\alpha = 1$ and $\beta = \frac{1}{4}$. Comparing Figures 6 and 7, it is clear that there is only a minor change in the equilibrium. As expected from (7'), there is now an additional support region S' , a narrow band that appears between the regions for S and NS (or $NS-L$). Notwithstanding this new region, the overall solution is still qualitatively similar to that in Figure 6. The equilibrium profit is again non-monotonic—as before, it is decreasing in c in regions where the vendor offers support, increasing in the limit region, and flat afterwards—implying that our earlier results with respect to the value of commitment are robust even to situations where consumers lack perfect information and differ in their beliefs of $\frac{c}{\gamma}$.

7.2 Partial Extension of Support

Our analysis has so far assumed that the time epochs in Figure 1 are all fixed and are exogenous in the model. An immediate implication is that, if support is ever extended beyond t_1 , it can only be extended all the way up to t_2 , that is, till the very end of the remaining useful life of the product. Thus, the support cost, c , which depends on the expected number of vulnerabilities patched during $[t_1, t_2]$, is proportional to $(t_2 - t_1)$ and is also fixed. How critically do our results depend on this

⁹The profit function for S'' is cubic in p . The first order condition yields two roots, of which only the smaller one satisfies the second order condition and yields the optimal price for S'' . It can be shown that this optimal price, $\frac{2c + \gamma(1+\mu) - \sqrt{4c^2 - 2c\gamma(1+\mu) + \gamma^2(1+\mu)^2}}{3\gamma}$, is increasing in μ and achieves its maximum of $\frac{c}{\gamma}$ when $\mu \rightarrow \infty$. This means that, for any finite μ , the optimal price must be below $\frac{c}{\gamma}$. However, that is a clear violation of $p > \frac{c}{\gamma}$, the condition that must be met for S'' .




Figure 7: Equilibrium Profit When Consumers Have Imperfect Information; $\alpha=1$, $\beta=\frac{1}{4}$

setting? What would happen if the vendor provides partial support, that is, support for only a fraction of the remaining useful life? To answer these questions, we now allow the vendor to provide fractional support of λ , $\lambda \in [0, 1]$, up to another time point $t'_2 = t_1 + \lambda(t_2 - t_1) \leq t_2$. Doing so would invariably mean proportional shrinkage in patch development by a factor of λ and thus a lower cost to extend support.

Note that, if $\lambda = 1$, the vendor's strategy coincides with S , and if $\lambda = 0$, it coincides with NS . In other words, a fractional λ corresponds to a setting that is somewhere between S and NS . Accordingly, we can write the total cost, patch development costs plus financial losses faced from leaving consumers partially unsupported, as:

$$\underbrace{\lambda c}_{\text{cost of partial support}} + \underbrace{(1 - \lambda)(\alpha + \beta\mu)\bar{v}}_{\text{losses from leaving consumers partially unprotected}} = \lambda c + \gamma\bar{v}(1 - \lambda)(1 + \mu). \quad (12)$$

That the patch development cost is linear in λ is fairly reasonable, since prior literature commonly assumes a steady rate for identification of vulnerabilities (Cavusoglu et al. 2008, Dey et al. 2015). At the same time, it is also logical that consumers' resentment, as well as the resulting fallout, is proportional to the length of the period for which they are denied necessary support. In fact, security-related losses faced by consumers should also scale proportionally with the length of the period without support, which leads us to:

$$\bar{v}(\lambda) = \frac{p}{1 + (1 - \lambda)\mu}. \quad (13)$$

It is apparent that (13) is a straightforward generalization of (1), from a discrete setup to a continuous one where fractional values of λ are also allowed along with the extremes of 0 and 1. We can now generalize (2) in Section 3 by combining (12) and (13), and express the profit as:

$$\pi(\lambda) = p \left(1 - \frac{p}{1 + (1 - \lambda)\mu} \right) - \left(\lambda c + (1 - \lambda) \frac{\gamma p (1 + \mu)}{1 + (1 - \lambda)\mu} \right). \quad (14)$$

The first order condition, $\frac{\partial \pi(\lambda)}{\partial \lambda} = 0$, is a quadratic equation in λ , which has at most one root less than one: $\lambda^* = \frac{1 + \mu}{\mu} - \frac{\sqrt{cp(\gamma(1 + \mu) - p\mu)}}{c}$. This λ^* is real only if $\gamma(1 + \mu) - p\mu > 0$; however, then, the second order condition is clearly violated:

$$\frac{\partial^2 \pi(\lambda)}{\partial \lambda^2} = \frac{2p\mu(\gamma(1 + \mu) - p\mu)}{(1 + (1 - \lambda)\mu)^3} > 0.$$

In other words, irrespective of the choice of p , an interior maximum simply does not exist. Therefore, the solution to the two-stage problem would only yield a boundary solution, $\lambda = 0$ or $\lambda = 1$, implying that the characterization in (4) is still valid.

Moving on to the issue of time-consistency, after consumers have made their upgrade decisions, \bar{v} is known, and the vendor seeks to choose a λ that minimizes $\lambda c + \gamma \bar{v}(1 - \lambda)(1 + \mu)$. Since this expression is linear in λ for a given \bar{v} , the vendor again chooses either $\lambda = 0$ or $\lambda = 1$. Naturally, the vendor would renege on its promise to end support and set $\lambda = 1$ when $c < \gamma \bar{v}(1 + \mu)$. Indeed, by the time support for Windows XP truly ended, more than 90% of users had switched to later

versions,¹⁰ making $\lambda = 1$ for all practical purposes. The implication is clear: Consistency issues persist exactly the same way, for the exactly same range of values of c identified in Proposition 1. Further, the no-commitment equilibrium is also identical, since the vendor is certain to pick an extreme value of λ when it moves after consumers' upgrade decisions. Thus, all our earlier results continue to hold as originally stated.

7.3 Impact of Normalization

In Sections 3 and 4, we worked with a normalized model in which consumers' valuation and the size of their base were both normalized. Simultaneous normalization of multiple parameters in a positive model can often be tricky, because spurious solutions can creep in at times. In this section, we wish to examine whether such normalization is without any loss of generality and if it impacts our results. To this end, we develop a non-normalized setup here and check if it is indeed equivalent to the normalized setup in the main paper. To avoid confusion, we employ slightly different notations here; we use \tilde{z} to denote the non-normalized counterpart of z , $z \in \{p, \alpha, \beta, \mu, c, \pi\}$.

Let us now consider the two parameters that we implicitly normalized. First, we assumed v to be uniformly distributed over $[0, 1]$. Second, we assumed a market size of one. To get rid of these normalizations, we now assume that v is uniformly distributed over $[0, M]$; that is, the maximum valuation is M . In a similar vein, we now let the market size to be N . Clearly, in the original model, M and N were both set to one.

The model formulation in the presence of these two additional parameters, M and N , is slightly different. This is because \bar{v} , which stands for the marginal user, is no longer the number of consumers sticking to the old version. More specifically, if \bar{n} is the number of such consumers, the following holds:

$$\frac{\bar{n}}{N} = \frac{\bar{v}}{M} \quad \Rightarrow \quad \bar{n} = \frac{N\bar{v}}{M}. \quad (15)$$

Equation (15) implies that (IR-NS) must now change to:

$$\bar{v} + \tilde{\mu} \frac{N\bar{v}}{M} \geq \tilde{p}, \quad (\text{IR-NS}'')$$

¹⁰A small number of users are still able to use XP using backdoors that make the operating system appear as Windows Embedded System, which will be supported till 2019. Also, some corporations convinced Microsoft to provide XP support for a price.

where \tilde{p} is the non-normalized price. Since (IR-S) is still valid, we ought to have:

$$\bar{v} = \begin{cases} \tilde{p}, & \text{for } S, \\ \frac{\tilde{p}}{1 + \frac{N}{M}\tilde{\mu}}, & \text{for } NS. \end{cases} \quad (1'')$$

The cost of extending support is fixed as before, except that we now denote it as \tilde{c} . Likewise, the potential losses that can result from not extending support is denoted $(\tilde{\alpha} + \tilde{\beta}\tilde{\mu})\bar{n}$, where \bar{n} is the number of unsupported users as given by (15). Note that, for *NS*:

$$(\tilde{\alpha} + \tilde{\beta}\tilde{\mu})\bar{n} = (\tilde{\alpha} + \tilde{\beta}\tilde{\mu})\frac{N\bar{v}}{M} = (\tilde{\alpha} + \tilde{\beta}\tilde{\mu})\frac{N\tilde{p}}{M(1 + \frac{N}{M}\tilde{\mu})} = MN \left(\frac{\tilde{\alpha}}{M} + \frac{\tilde{\beta}}{N} \frac{N}{M}\tilde{\mu} \right) \frac{\tilde{p}}{M(1 + \frac{N}{M}\tilde{\mu})}.$$

We are now ready to restate the vendor's profit maximization problem originally specified in (2) in Section 3:

$$\tilde{\pi} = \begin{cases} \tilde{p} \left(N - \frac{N\tilde{p}}{M} \right) - \tilde{c}, & \text{for } S, \\ \tilde{p} \left(N - \frac{N\tilde{p}}{M(1 + \frac{N}{M}\tilde{\mu})} \right) - MN \left(\frac{\tilde{\alpha}}{M} + \frac{\tilde{\beta}}{N} \frac{N}{M}\tilde{\mu} \right) \frac{\tilde{p}}{M(1 + \frac{N}{M}\tilde{\mu})}, & \text{for } NS. \end{cases}$$

Let us now define the following:

$$p = \frac{\tilde{p}}{M}, \quad \alpha = \frac{\tilde{\alpha}}{M}, \quad \beta = \frac{\tilde{\beta}}{N}, \quad \mu = \frac{N\tilde{\mu}}{M}, \quad \text{and } c = \frac{\tilde{c}}{MN}, \quad (16)$$

and substitute accordingly to obtain:

$$\tilde{\pi} = \begin{cases} MN(p(1-p) - c), & \text{for } S, \\ MN \left(p \left(1 - \frac{p}{1+\mu} \right) - \gamma p \right), & \text{for } NS, \end{cases} \quad (2'')$$

where $\gamma = \frac{\alpha + \beta\mu}{1 + \mu}$ as before. Comparing (2'') with (2), it is immediate that $\tilde{\pi}$ is simply MN times π .

To see the impact on our analysis in Section 4, we need to derive the counterpart of condition (6); recall that the primary difference between the vendor's problems in Sections 3 and 4 is precisely this condition. As before, we need to derive it by comparing the costs of extending support with

potential losses from not extending:

$$(\tilde{\alpha} + \tilde{\beta}\tilde{\mu}) \frac{N\tilde{p}}{M(1 + \frac{N}{M}\tilde{\mu})} \leq \tilde{c}.$$

Dividing both sides by MN and making the substitutions defined in (16), we get $p \leq \frac{c}{\gamma}$, which is exactly the same as the condition in (6). Thus, (6) is still applicable, and we can write the vendor's profit as follows:

$$\tilde{\pi} = \begin{cases} MN(p(1-p) - c), & \text{if } p > \frac{c}{\gamma} \text{ (S)}, \\ MN\left(p\left(1 - \frac{p}{1+\mu}\right) - \gamma p\right), & \text{if } p \leq \frac{c}{\gamma} \text{ (NS)}. \end{cases} \quad (7'')$$

Again, the profit in (7'') is MN times that in (7).

It should now be obvious that the normalization used in the main paper has no impact on our findings; it simply re-scales the decision variable and parameters according to (16) and the objective function by a factor of MN . Since a scaling of the objective function has no impact on various thresholds discussed in the main paper, the insights there should remain applicable even when all normalizations are done away with.

8 Conclusion

How many times do you have to pronounce the same software product dead? If recent developments are any indication, the answer is numerous times. Often, it is easier to pronounce the end of licensing life for an aging version but much harder to convince users to move to a new one. Furthermore, if a sizable section of the consumer base clings on to the old version, it turns out to be quite difficult for the software vendor to actually end all support as pronounced, as doing so runs the risk of significant security breaches and consequent losses in reputation. This quandary and the commitment problem that results from it are at the focus of this research.

We analyze the vendor's decision problem with respect to the optimal support strategy and the price for the upgrade. While doing so, we account for the cost of extending support, the security risks faced by consumers, and the potential loss the vendor faces as a result of consumer disenchantment. We begin by analyzing the game in which the vendor first announces its decisions

and consumers subsequently make their choices. We show that, in this simple setup, the vendor’s optimal decision is time-inconsistent. In other words, the vendor has an incentive to renege on its plan to end support after consumers have made their upgrade decisions. Accordingly, we analyze the no-commitment equilibrium, in which consumers act as if the vendor is not committed to a support strategy.

We find that the commitment problem faced by the vendor adversely affects its profit. The apparent flexibility to readjust the support strategy does not help the vendor. The flexibility is, in fact, detrimental to its profit. This is because, when consumers suspect that the vendor may backtrack from its promise to end support, they can no longer count on the security-related benefits of upgrading. As a result, their willingness to pay for the new version becomes depressed, which, in turn, limits the vendor’s ability to profit from the newly launched version.

The commitment problem has significant implications for the vendor’s optimal pricing strategy. In the no-commitment equilibrium, the vendor ought to consider options beyond the usual support and no-support strategies. Interestingly, as we show in Figure 5, when the cost to support is moderate, it might become imperative for the vendor to pursue a “limit pricing” strategy to signal a firm commitment to end support. By holding the price down to a limiting value, p_L , the vendor is able to send the message that the lower price would lead to only a handful of users clinging to the older version, effectively eliminating the need to further extend support. If the vendor fails to recognize the importance of using such a price reduction to signal a firm commitment, and mistakenly sets the price high, the consequences can be serious—if consumers respond strategically, not only would the vendor end up extending support but it would also lose significant amounts in profits. In short, a reduced price for the new version, despite its apparent unattractiveness, could actually help the vendor by credibly signaling the intent to end critical support.

Counterintuitively, in the no-commitment equilibrium, the vendor’s profit may be non-monotonic in the cost of support, implying that it might help the vendor to look for ways to inflate that cost, instead of simply choosing the optimal strategy for a given cost; see Figure 6. This is contrary to the common wisdom that higher costs are necessarily unwelcome. In essence, a higher cost in this context simply serves as a credible signal that the vendor would not extend support, and eliminates doubts in consumers’ minds about the security-related benefits of the upgrade, which, in turn, allows the vendor to increase the price of the new version.

The non-monotonicity of profit also provides us important cues on designing a suitable commitment mechanism. As Proposition 4 shows, when faced with the commitment problem, the vendor can do better by offering a contingent refund, payable to buyers of the new version in the event support is extended for the old one. Although similar refunds have been offered in other contexts, the notion of using them as a means to effectively commit to ending support is novel here, and it does provide a software vendor with a practical way to address the commitment problem faced when announcing the demise of an aging software.

The welfare implications are also intriguing. Intuitively, one would expect continued support to be beneficial for consumers and the society. We find that this is not necessarily the case. In fact, even if the vendor ends support, consumers are not necessarily worse off. As Proposition 5 shows, when security risks and the resulting network effects are not high compared to the vendor's loss rate, the vendor actually needs to be "cruel to be kind" (Kelly 2014). In this situation, facing a somewhat high loss rate, the vendor's natural tendency is to minimize the number of consumers who would ultimately be without necessary security support. Accordingly, it sets a lower upgrade price to entice more consumers to move to the new version. And, as we discover, the gain from more consumers using a better product can indeed dominate the drawback of leaving a few without support. This finding is of considerable interest, since some experts and government officials often equate ending support of a popular product to something bad for consumers as well as the society as a whole. We show that such thinking and any myopic arm-twisting of the vendor for extended support are clearly not justified across the board. Not only consumer welfare can get compromised as a result of such actions, social welfare may suffer as well.

Although this study reveals several useful insights, it is not without limitations. Further research is necessary for a more refined understanding of the vendor's strategic options. In this research, we have considered one old product and one new. However, in a few years, the new product would become old, and a newer product would enter the market, implying that the game might repeat itself. This is fairly significant, since "the outcome is vastly different when there is a future" (Zupan 2011, p.194). Consider, for example, a game between an incumbent and an entrant firm where the entrant is unaware of the "type"—weak or strong—of the incumbent (Fudenberg and Tirole 1991, p.371). The incumbent obviously wants to enjoy a monopoly, but if it is weak, fighting back is too costly from a near-term perspective. Yet, one of the equilibriums is such that even a weak

incumbent fights back the first entrant just to appear strong to the future ones.

In our context as well, repetition means an opportunity to influence the future. More specifically, pursuing self-interest in a repeated setting is quite a bit different from maximizing the profit function considered in this paper. Acting in self-interest would also involve taking into consideration informational and reputational effects of renegeing from a pre-announced strategy. For example, renegeing might help the current-period profit, but doing so can make consumers more likely to believe that the firm would do it again. Conversely, not renegeing could be costly in the near term but can serve as a signal to consumers that the vendor would follow through on its future announcements. Currently, our model does not capture these factors and ignores the dynamics of a repeated setting entirely. We hope to address this limitation in upcoming endeavors for a more complete understanding of this context.

References

- Arora, A., J.P. Caulkins, R. Telang. 2006. Sell first, fix later: Impact of patching on software quality. *Management Science* **52**(3) 465–471.
- Arora, A., R. Krishnan, R. Telang, Y. Yang. 2010. An empirical analysis of software vendors' patch release behavior: Impact of vulnerability disclosure. *Information Systems Research* **21**(1) 115–132.
- Arora, A., R. Telang, H. Xu. 2008. Optimal policy for software vulnerability disclosure. *Management Science* **54**(4) 642–656.
- August, T., M.F. Niculescu, H. Shin. 2014. Cloud implications on software network structure and security risks. *Information Systems Research* **25**(3) 489–510.
- August, T., T.I. Tunca. 2006. Network software security and user interface. *Management Science* **52**(11) 1703–1720.
- August, T., T.I. Tunca. 2008. Let the pirates patch? an economic analysis of network software security patch restrictions. *Information Systems Research* **19**(1) 48–70.
- August, T., T.I. Tunca. 2011. Who should be responsible for software security? a comparative analysis of liability policies in network environments. *Management Science* **57**(5) 934–959.
- Bala, R., S. Carr. 2009. Pricing software upgrades: The role of product improvement and user costs. *Production and Operations Management* **18**(5) 560–580.
- Bhattacharya, S., V. Krishnan, V. Mahajan. 2003. Operationalizing technology improvements in product development decision-making. *European Journal of Operational Research* **149**(1) 102–130.
- Cavusoglu, H., H. Cavusoglu, J. Zhang. 2008. Security patch management: Share the burden or share the damage? *Management Science* **54**(4) 657–670.
- Choudhary, V., Z. Zhang. 2015. Patching the cloud: The impact of SaaS on patching strategy and the timing of software release. *Information Systems Research* **26**(4) 845–858.

- Dey, D., A. Lahiri, G. Zhang. 2012. Hacker behavior, network effects, and the security software market. *Journal of Management Information Systems* **29**(2) 77–108.
- Dey, D., A. Lahiri, G. Zhang. 2014. Quality competition and market segmentation in the security software market. *MIS Quarterly* **38**(2) 589–606.
- Dey, D., A. Lahiri, G. Zhang. 2015. Optimal policies for security patch management. *INFORMS Journal on Computing* **27**(3) 462–477.
- Dhebar, A. 1994. Durable-goods monopolists, rational consumers, and improving products. *Marketing Science* **13**(1) 100–120.
- Ellison, G., D. Fudenberg. 2000. The neo-luddite’s lament: Excessive upgrades in the software industry. *The RAND Journal of Economics* **31**(2) 253–272.
- Evangelista, B. 2014. Microsoft to end Windows XP support April 8. <http://www.sfgate.com/technology/article/Microsoft-to-end-Windows-XP-support-April-8-5298608.php>. Accessed on May 21, 2015.
- Fudenberg, D., J. Tirole. 1991. *Game Theory*. 1st ed. MIT Press, Cambridge, MA.
- Galbreth, M.R., M. Shor. 2010. The impact of malicious agents on the enterprise software industry. *MIS Quarterly* **34**(3) 595–612.
- Ioannidis, C., D. Pym, J. Williams. 2012. Information security trade-offs and optimal patching policies. *European Journal of Operational Research* **216**(2) 434–444.
- Kannan, K., M. Rahman, M. Tawarmalani. 2016. Economic and policy implications of restricted patch distribution. *Management Science* **62**(11) 3161–3182.
- Kannan, K., R. Telang. 2005. Market for software vulnerabilities? Think again. *Management Science* **51**(5) 726–740.
- Keizer, G. 2012. Adobe: Pay upgrade price to patch critical bugs. <http://www.computerworld.com/article/2504314/malware-vulnerabilities/adobe-pay-upgrade-price-to-patch-critical-bugs.html>. Accessed on May 21, 2015.
- Keizer, G. 2014. Perspective: Microsoft risks security reputation ruin by retiring XP. <http://www.computerworld.com/article/2488375/microsoft-windows/perspective-microsoft-risks-security-reputation-ruin-by-retiring-xp.html>. Accessed on May 18, 2015.
- Kelly, G. 2014. Microsoft saves Windows XP in an act of utter stupidity. <http://www.forbes.com/sites/gordonkelly/2014/05/02/microsoft-saves-windows-xp-in-an-act-of-utter-stupidity/>. Accessed on June 10, 2015.
- Kim, B.C., P.-Y. Chen, T. Mukhopadhyay. 2011. The effect of liability and patch release on software security: The monopoly case. *Production and Operations Management* **20**(4) 603–617.
- Kornish, L.J. 2001. Pricing for a durable-goods monopolist under rapid sequential innovation. *Management Science* **47**(11) 1552–1561.
- Kydland, F.E., E.C. Prescott. 1977. Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy* **85**(3) 473–492.
- Lahiri, A. 2012. Revisiting the incentive to tolerate illegal distribution of software products. *Decision Support Systems* **53**(2) 357–367.
- Mankiw, N.G. 1988. Recent developments in macroeconomics: A very quick refresher course. *Journal of Money, Credit and Banking* **20**(3) 436–449.
- Mehra, A., R. Bala, R. Sankaranarayanan. 2012. Competitive behavior-based price discrimination for software upgrades. *Information Systems Research* **23**(1) 60–64.

- Mehra, A., A. Seidmann, P. Mojumder. 2014. Product life-cycle management of packaged software. *Production and Operations Management* **23**(3) 366–378.
- Osborne, Martin J. 2004. *An Introduction to Game Theory*. 1st ed. Oxford University Press, New York, NY.
- Padmanabhan, V., S. Rajiv, K. Srinivasan. 1997. New products, upgrades, and new releases: A rationale for sequential product introduction. *Journal of Marketing Research* **34**(4) 456–472.
- Png, I.P.L., Q.-H. Wang. 2009. Information security: Facilitating user precautions vis-à-vis enforcement against attackers. *Journal of Management Information Systems* **26**(2) 97–121.
- Protalinski, E. 2014. Microsoft offers prorated Office 365 refunds to paid subscribers after making mobile editing free. <http://venturebeat.com/2014/11/07/microsoft-offers-pro-rated-office-365-refunds-to-paid-subscribers-after-making-mobile-editing-free/>. Accessed on May 21, 2015.
- Ramachandran, K., V. Krishnan. 2008. Design architecture and introduction timing for rapidly improving industrial products. *Manufacturing & Service Operations Management* **10**(1) 149–171.
- Ransbotham, S., S. Mitra. 2009. Choice and chance: A conceptual model of paths to information security compromise. *Information Systems Research* **20**(1) 121–139.
- Reuters. 2014. China bans use of Microsoft’s Windows 8 on government computers. <http://www.reuters.com/article/2014/05/20/us-microsoft-china-idUSBREA4J07Q20140520>. Accessed on May 18, 2015.
- Sankaranarayanan, R. 2007. Innovation and the durable goods monopolist: The optimality of frequent new-version releases. *Marketing Science* **26**(6) 774–791.
- Seltzer, L. 2014. Microsoft to extend Windows XP anti-malware updates one year. <http://www.zdnet.com/article/microsoft-to-extend-windows-xp-anti-malware-updates-one-year/>. Accessed on May 18, 2015.
- Sun, L. 2014. The real reason China banned Microsoft’s Windows. <http://www.fool.com/investing/general/2014/05/22/the-real-reason-china-banned-microsofts-windows-8.aspx>. Accessed on September 20, 2015.
- Timberg, C., E. Nakashima. 2014. Government computers running Windows XP will be vulnerable to hackers after April 8. http://www.washingtonpost.com/business/technology/government-computers-running-windows-xp-will-be-vulnerable-to-hackers-after-april-8/2014/03/16/9a9c8c7c-a553-11e3-a5fa-55f0c77bf39c_story.html. Accessed on May 21, 2015.
- Zupan, M.A. 2011. The virtues of free markets. *Cato Journal* **26**(2) 171–198.

Appendix: Proofs

Proof of Proposition 1

From (2), we get:

$$\frac{d\pi}{dp} = \begin{cases} 1 - 2p, & \text{for } S, \\ 1 - \gamma - \frac{2p}{1+\mu}, & \text{for } NS. \end{cases}$$

It is immediate that the optimal price and profit are indeed as given by (3). It is also easy to verify that both p_S and p_{NS} in (3) satisfy the required second order conditions. Further, solving $\pi_S = \pi_{NS}$, where π_S and π_{NS} are as in (3), leads to the threshold $c = c_1$; above this threshold, NS becomes the dominant strategy as indicated in (4).

Now, after consumers have made their upgrade decisions, the vendor reinstates support if doing so is relatively less costly when compared to financial losses likely to stem from ending support, that is, if

$$c < (\alpha + \beta\mu)\bar{v}.$$

When $c > c_1$ and the vendor has chosen *NS*, we ought to have $\bar{v} = \frac{p}{1+\mu}$; see (1). Substituting for \bar{v} and recalling the definition of γ from Section 3.3, we can reduce the above condition to:

$$\frac{c}{\gamma} < p.$$

Substituting p_{NS} for p , we can further reduce this condition to:

$$c < c_2,$$

where c_2 is as defined in the statement of this proposition. Thus, when $c_1 < c < c_2$, *NS* is optimal but not time-consistent.

Finally, as illustrated using numerical examples in the main paper, $c_1 < c_2$ is very much a possibility when γ is relatively low. Therefore, time-consistency of the solution in (4) cannot be guaranteed. ■

Proof of Proposition 2

The prices and profits for *S* and *NS* are still as given by (3), respectively. For *NS-L*, we have:

$$p^* = p_L = \frac{c}{\gamma} \quad \text{and} \quad \pi^* = \pi_L = c \left(\frac{1}{\gamma} \left(1 - \frac{c}{\gamma(1+\mu)} \right) - 1 \right).$$

As explained in the proof of Proposition 1, $c < c_2$ is equivalent to $p_{NS} > \frac{c}{\gamma}$, which violates the condition for *NS* in (7). So, when $c < c_2$, we only need to compare *NS-L* with *S* to see which one is more profitable. Setting $\pi_S = \pi_L$, we get the desired threshold:

$$c_3 = \frac{\gamma(1+\mu)}{2} \left(1 - \sqrt{\frac{\mu}{1+\mu}} \right),$$

with $\pi_L > \pi_S$ for $c > c_3$. On the other hand, if $c \geq c_2$, *NS* is possible. Since the interior solution must dominate any other solution, $\pi_{NS} \geq \pi_L$ holds for $c \geq c_2$, with the inequality becoming strict for $c > c_2$. Therefore, *NS-L* is the optimal choice only when $c_3 < c < c_2$.

From Proposition 1, we know that $\pi_{NS} > \pi_S$ when $c > c_1$. This, taken together with the discussion in the preceding paragraph, implies that *NS* is optimal when $c \geq c_2$ as well as $c > c_1$. Now,

$$\frac{c_1 - c_2}{c_3 - c_2} = \frac{1}{2} \left(1 + \frac{1}{\gamma} \sqrt{\frac{\mu}{1+\mu}} \right) > 0,$$

implying that $c_1 > c_2$ if and only if $c_3 > c_2$. So, when $c_3 \leq c_2$, $c \geq c_2$ is sufficient since it also implies that c must be at least as large as c_1 . Likewise, if $c_3 > c_2$, $c_1 > c_2$ is automatically satisfied, and $c > c_1$ is sufficient to guarantee *NS* is optimal.

Having delineated the region of optimality for *NS-L* and subsequently for *NS*, it is easy to do the same for *S*, as it has to be the preferred strategy in the portion of the parameter space where the other two are

not optimal. Further, since

$$\frac{c_3}{\gamma} = \frac{1+\mu}{2} \left(1 - \sqrt{\frac{\mu}{1+\mu}}\right) < \frac{1}{2}, \quad \forall \mu > 0,$$

it is immediate that $p_S > \frac{c}{\gamma}$ holds for $c \leq c_3$. Now, note that $c \leq c_3$ must hold for S to be optimal, since the vendor would choose $NS-L$ otherwise. Therefore, $p_S > \frac{c}{\gamma}$ is satisfied throughout the region of optimality of S , just as required by (7). ■

Proof of Proposition 3

Recall from (4) that $\pi_{NS} > \pi_S$ when $c > c_1$. Further, since the interior solution must dominate any other solution, $\pi_{NS} \geq \pi_L$ holds everywhere, with the inequality becoming strict for $c \neq c_2$. Therefore, when $c_1 < c < c_2$, π_{NS} is higher than both π_S and π_L and, hence, the no-commitment equilibrium profit in (9).

It is also immediate from (9) that

$$\frac{d\pi_S}{dc} = -1, \quad \text{and} \quad \frac{d\pi_{NS}}{dc} = 0,$$

implying that the profit is decreasing in c in Region S but independent of c in Region NS . To show that the profit is increasing in c in Region $NS-L$, note that

$$\frac{d^2\pi_L}{dc^2} = -\frac{2}{\gamma^2(1+\mu)},$$

which is negative. Therefore, $\frac{d\pi_L}{dc}$ attains its minimum at the upper boundary of Region $NS-L$, which is $c = c_2$. However, the minimum value itself is positive, since

$$\left. \frac{d\pi_L}{dc} \right|_{c=c_2} = \mu > 0.$$

So, $\frac{d\pi_L}{dc}$ must be positive everywhere in Region $NS-L$. ■

Proof of Proposition 4

Let the desired refund be r per consumer. This refund imposes an extra cost burden on the vendor, so it would change the condition in (6) to:

$$c + r \left(1 - \frac{p}{1+\mu}\right) \geq \gamma p.$$

Setting $p = p_{NS}$, we can reduce the above inequality to:

$$r \geq \frac{\gamma(1-\gamma)(1+\mu) - 2c}{1+\gamma}.$$

Therefore, the required minimum refund amount is $\rho = \frac{\gamma(1-\gamma)(1+\mu) - 2c}{1+\gamma}$. ■

Proof of Proposition 5

From (11), we can see that CS is (weakly) higher for strategy NS when compared to S when:

$$\phi_{NS} = \frac{(1 + \gamma)^2 - 4\mu(1 - \gamma)}{8} > \frac{1}{8} = \phi_S.$$

Rearranging the terms, we obtain the condition in the theorem, which is:

$$\gamma > \left(\sqrt{1 + 4\mu(2 + \mu)} - (1 + 2\mu) \right) = \Theta(\mu).$$

Clearly, for a given γ and μ , consumers prefer NS to S under this condition.

The remaining question is whether $NS-L$ is also preferable to S . To answer this, we note that, according to (11), the consumer surplus for NS does not depend on c anywhere, and it also matches that for $NS-L$ at the boundary $c = c_2$. However, for $NS-L$, the consumer surplus does depend on c , and its derivative with respect to c is $\frac{c - \gamma(1 + \mu)^2}{\gamma^2(1 + \mu)^2}$. Now, since $c \leq c_2$, we can write:

$$\frac{c - \gamma(1 + \mu)^2}{\gamma^2(1 + \mu)^2} \leq \frac{c_2 - \gamma(1 + \mu)^2}{\gamma^2(1 + \mu)^2} = -\frac{(1 + \gamma + 2\mu)}{2\gamma(1 + \mu)} < 0.$$

Because the derivative is always negative within Region $NS-L$, the consumer surplus must be even higher than what NS can possibly yield. Thus, if the condition $\gamma > \Theta(\mu)$ holds, consumers would not only prefer NS to S but would also prefer $NS-L$ to S .

When $\gamma \leq \Theta(\mu)$, consumers clearly prefer S to NS . However, they would prefer $NS-L$ to S as long as

$$\phi_L = \frac{c^2 - \gamma(1 + \mu)^2(2c - \gamma)}{2\gamma^2(1 + \mu)^2} > \frac{1}{8} = \phi_S.$$

Upon simplification, this leads to

$$\gamma > \frac{2c \left(2(1 + \mu) + \sqrt{1 + 4\mu(2 + \mu)} \right)}{3(1 + \mu)} = \theta(c, \mu).$$

Note that the expression on the right is increasing in c , and at $c = c_2$, it does coincide with $\Theta(\mu)$. ■