Cloud Implications on Software Network Structure and Security Risks

Terrence August[*]
Rady School of Management
University of California, San Diego

Marius Florin Niculescu[†]
College of Management
Georgia Institute of Technology

Hyoduk Shin[‡]
Rady School of Management
University of California, San Diego

March 2013

Abstract

By software vendors offering, via the cloud, *software as a service* (SaaS) versions of traditionally on-premises products, security risks associated with software usage become more diversified which can greatly increase the value associated with network software. In an environment where negative security externalities are present and users make complex consumption and patching decisions, we construct a model that clarifies whether and how SaaS versions should be offered by vendors. For high security-loss software, we find that SaaS should be geared to the middle tier of the consumer market if patching costs and the quality of the service offering are high, and geared to the lower tier otherwise. The vendor-preferred and social planner-preferred structure of segmentation often coincide, except when patching costs are within an intermediate range, in which case welfare can be increased by additional incentives to induce SaaS usage in the middle tier. Relative to on-premises benchmarks, we find that software diversification indeed leads to lower average security losses for users when patching costs are high. However, when patching costs are low, surprisingly, average security losses can actually increase as a result of SaaS offerings and lead to lower consumer surplus. Releasing a SaaS version substantially increases vendor profits and welfare in high security-loss environments and only has a limited, but positive, effect in low security-loss environments. Nevertheless, in the latter case, we find that versioning is optimal in the presence of version-specific negative security externalities, even as they become small.

[*] Rady School of Management, University of California, San Diego, La Jolla, CA 92093-0553. e-mail: taugust@ucsd.edu

[†] College of Management, Georgia Institute of Technology, Atlanta, GA 30308. e-mail: marius.niculescu@mgt.gatech.edu

[‡] Rady School of Management, University of California, San Diego, La Jolla, CA 92093-0553. e-mail: hdshin@ucsd.edu

1 Introduction

With broadband access becoming faster and more pervasive, there has been a shift back toward software deployment models where computing is centralized and accessed via thin clients. Cloud computing has emerged both out of economic efficiency and to satisfy users' preferences for their data and applications to be ubiquitous. Because interaction with the cloud is now much faster, both consumers and businesses need not rely as much on personal computing power. Over the last two decades, consumers have increasingly harnessed the cloud for personal e-mail, online gaming, photo sharing, and social networking; and, more recently, they are considering the cloud for provision of productivity software such as Google Docs and Microsoft Office 365 (El Akkad 2011). Similarly, over time businesses have moved toward creating web applications to support business processes which utilize web browsers as a natural front-end easily accessible by employees. In this case, computation shifted from heavyweight clients to back-end servers maintained internally by IT staff. However, with more diverse cloud offerings such as Amazon's Elastic Compute Cloud, Salesforce CRM solutions, and Google Apps for Business, even the internal back end systems are no longer a necessity and can be transitioned to the cloud.

Many governments are also starting to implement cloud-based models to support and make business and government processes more operationally efficient. For example, the U.S. government which has an $80 billion federal IT budget, has championed a Federal Cloud Computing Initiative to encourage agencies to move toward cloud computing solutions, supporting this transition with Apps.gov (Claburn 2009). Gartner estimates that the cloud computing industry will grow to $149 billion by 2015 (Kundra 2011). Vivek Kundra, chief information officer of the U.S., also suggests cloud computing will help increase productivity in health care, financial services, and education, pointing out that a one percent productivity increase in health care over the next 10 years represents $300 billion in value (e.g., shifting electronic medical records to the cloud). However, cloud computing is not likely to be an "end all" solution. Rather, for many firms, cloud computing will be a single component of an overall IT strategy that augments the traditional use models currently employed (O'Neill 2011).

As more users are willing to use *software as a service* (SaaS - applications ran from the cloud), vendors will readily begin developing these offerings for many classes of software applications. From a vendor's perspective, SaaS has many benefits when compared to its *on-premises* counterpart including significantly lower piracy, reduced distribution costs, and, particularly, greater control over security. For example, with on-premises software, it is difficult to incentivize users to patch their installations when security patches are released, and poor user patching behavior reduces the value of the software vendor's product. For on-premises offerings such as Microsoft Windows, Oracle, and Microsoft IIS, the user network is characterized by a large number of widespread nodes where individual instances of the software are running with many remaining unpatched (Lemos 2004, Keizer 2008). Large installed software networks such as these are primary targets for *undirected* attacks via computer worms and viruses; Code Red, SQL Slammer, Sasser, and Conficker are all examples of malware that spread across vulnerable software networks and caused sizable economic damages (Moore et al. 2002, Lemos 2003, Keizer 2004, Markoff 2009).

On the other hand, if a vendor releases a SaaS version of its software product, it can ensure its hosted software always has the latest security patches applied. SaaS offerings tend to be centralized on the provider's servers and are less prone to these undirected attacks. Nevertheless, because of the magnitude of user information located in one place, SaaS may be more susceptible to *directed* attacks in which individuals with malicious intent specifically target and attack the vendor's systems. Recently, Google, Salesforce, and Sony have been victimized by directed attacks associated with their SaaS offerings. Google lost intellectual property after its systems were compromised by an exploit on a zero-day vulnerability in Internet Explorer, while senior U.S. officials and hundreds of others had their privacy breached (Gorman and Vascellaro 2010, Zetter 2010, Efrati and Gorman 2011). In 2007, a Salesforce.com employee fell victim to a phishing attack which led to significant exposure of customer information (Espiner 2007). Similarly, in a very recent attack, Sony's PlayStation network was penetrated leading to over 100 million user accounts being compromised (Sherr and Wingfield 2011). In addition to individuals' losses, estimates suggest Sony may lose more than $1.25 billion due to lost business, compensation costs, and investments, while banks may incur up to $300 million in credit card replacement costs (Aspan and Baldwin 2011, Osawa 2011).

Whether software is deployed in a traditional on-premises fashion or in the SaaS paradigm, security attacks on software vulnerabilities will continue to be a challenging problem. Yet, in terms of software security risk management, a trend toward increasing SaaS usage may actually help reduce total risk by diversifying exposure across undirected and directed attacks and limiting the sizes of particular populations that malware can effectively target; this, in turn, may indirectly reduce the incentives of malware developers to target diversified software (Bain et al. 2002).

In this paper, we formally examine how a SaaS offering by a software vendor affects the security risk faced by users of his software network and the total value derived from his software product. We study his product differentiation and pricing problem under distinct security externalities in each paradigm and examine how aggregate security risk on the network is affected. To do so, we build a model of consumer behavior that captures their incentives to use either the on-premises or SaaS version of software and, in the case of the former, to patch their individual installations when security vulnerabilities arise. Two important features in our model are that: (i) users of the on-premises alternative who choose not to patch cause negative security externalities on other users, and (ii) all users of the SaaS version cause a negative externality on other users of SaaS by increasing the aggregate likelihood of a directed attack due to increased valuable information stored at a centralized location.

Using our model, we first study how consumers behave in equilibrium in the face of security externalities and characterize the manner in which they segment across alternatives and patching strategies for varying security-loss environments. We then examine how a software vendor sets prices of his on-premises and SaaS versions to induce profitable usage behavior. In particular, for high security-loss environments, where consumers are at risk of taking large losses if struck by security attacks, we establish that the vendor will cater his SaaS offering to the middle tier of the consumer market but only when both patching costs and the quality of the SaaS alternative are high, thus efficiently splitting on-premises usage characterized by higher valuation patched users and lower valuation unpatched users. Otherwise, he targets his SaaS offering to the lower tier of the consumer market. Because of his incentives to target the lower tier in the latter case, we also demonstrate that social welfare can be improved if he is incentivized to gear it back toward the middle tier.

For low security-loss environments, we again establish that the software vendor sets price to induce usage of both versions which permits a comparison between our results under security externalities and those found in the information goods versioning literature. We show that as long as each of the distinct versions of a given software product has a small amount of idiosyncratic risk stemming from its own user population, then versioning is optimal.

Next, we turn our attention to clarifying the benefits of introducing SaaS alternatives by comparing measures of profitability, social welfare, security risk, and consumer surplus with those obtained under the traditional, on-premises only benchmark. We demonstrate that there are substantial benefits to profit and welfare in high security-loss environments associated with introducing a SaaS alternative, but they are limited in low security-loss ones. Surprisingly, we find that in high security-loss environments, although releasing SaaS often improves aggregate security, in some cases its release can actually increase average per-user security losses - particularly, when patching costs are not too high. Furthermore, consumer surplus can also decrease in such a case if the quality of the SaaS alternative is sufficiently high. Finally, in light of our findings on the merits of introducing SaaS, comparing its benefits to direct reductions in the likelihood of security attacks and the magnitude of patching costs, we demonstrate that the SaaS strategy can be outperformed by reductions in one of these parameters in both sufficiently low and sufficiently high security-loss environments. However, for more moderate ones, the magnitude of the direct cost reduction would need to be quite high, suggesting that the diversification benefits in a SaaS strategy would be preferable.

The rest of this paper is organized as follows. In Section 2, we review the related literature and, in Section 3, we formally present our model and characterize the consumer market equilibrium. In Section 4, we study the software vendor's pricing problem and clarify how he targets his SaaS offering. In Section 5, we provide comparisons with the benchmark case, and, in Section 6, we offer our concluding remarks. All proofs are provided in the Appendix.

2 Literature Review

Our primary contribution lies within the literature on the economics of information security where we focus on security considerations associated with a software vendor's decision to expand a software product's offering to include SaaS.[1] Our study is the first to examine software diversification in the presence of negative security externalities that users impose on one another with both their usage choice and patching behavior. To develop our work, we draw from several streams of literature which we describe below.

Our model is closest in spirit to prior studies that endogenize the consumer decision to patch in

[1] Anderson (2001), Gordon and Loeb (2002), Anderson and Moore (2006), and Johnson (2008) together provide a comprehensive introduction to important themes in this research area.

the face of security risks (e.g., August and Tunca 2006, Choi et al. 2010). In particular, we build upon the model in August and Tunca (2006) which captures how risk faced by unpatched users is related to the number of users who choose to be unpatched in equilibrium. Whereas the prior literature above has investigated how patching rebates, mandates, and taxes can improve software security, our focus in the current work is to study how introducing both on-premises and SaaS alternatives affects usage decisions across alternatives and how this behavior affects the security properties of the network.

This paper complements other research streams on piracy, software liability, and vulnerability disclosure, which, similar to the current work, all study particular facets of the security problem and recommend strategies to manage risk and improve the value derived from software. Several studies extend the consumer type space to account for piracy and then explore how piracy interacts with software security risk (see, e.g., Rahman et al. 2007, August and Tunca 2008, Lahiri 2011). Kim et al. (2010, 2011) explore vendor liability on software security losses and how risk sharing of security losses increases software quality. They establish that software liability helps increase security quality when consumers are heterogeneous in their sensitivities to quality. Studying loss liability on zero-day attacks, liability on patching costs, and security standards, August and Tunca (2011) find that security standards perform best in low zero-day risk environments, while patch liability is most effective in high ones. Notably, Cavusoglu et al. (2008) also study cost sharing and loss liability schemes in the context of optimal time-driven patch management, finding that cost-sharing performs better. In the current work, we abstract from specific timing issues related to vulnerability disclosure for which there is a well-developed literature (see, e.g., Cavusoglu et al. 2007, Arora et al. 2008, Choi et al. 2010, Ransbotham et al. 2010), taking security patches as being available at the time of patching as in August and Tunca (2006, 2008).

Our work is also topically related to some important studies that examine the connection between information security, diversification, and types of security attacks. In Chen et al. (2011), the authors are the first to construct a model that explores the trade-offs among increased security through software diversity, lost network effects, and economies of scale. They find that a firm can benefit more from diversification as software begins to use more standardized interfaces and when adapters and middleware are available to keep applications compatible. We complement their study by looking instead at the software vendor's decision to create diversified offerings (SaaS versus on-premises) in terms of usage structure in the presence of negative security externalities. Png and Wang (2009) examine the role of government in facilitating end-user precautions and enforcing laws against attackers. They find that facilitating end-user precautions tends to be more effective. Similar to their model, we also distinguish between directed and undirected attacks. In particular, we model how the diversified offerings vary in their susceptibility to each type of attack in order to analyze their aggregate impact on risk. Ransbotham and Mitra (2009) develop a model which clarifies the process that information security compromises follow. Stemming from their empirical analysis, they recommend utilization of vulnerability controls such as early patching to inhibit the progression of attacks. Dey et al. (2012) study the security software market and use a theoretical model to explain idiosyncratic characteristics of the market such as low market coverage yet high prices.

Lastly, our work is related to several papers that study various aspects of SaaS versus on-premises business models. Choudhary (2007) examines how SaaS versus perpetual licensing affect a software vendor's incentives to invest in quality. In a two-period model, he establishes that the vendor tends to invest more in quality under a SaaS scheme and that both profits and welfare increase as a result. Zhang and Seidmann (2010) study the licensing problem under network effects and quality uncertainty. They demonstrate that under strong network effects, hybrid models are favorable; in our work, we establish a similar result driven by security risk diversification benefits in contrast to multi-period dynamics. Huang and Sundararajan (2005) take a more general approach to pricing to characterize optimal non-linear prices of on-demand computing, while Ma and Seidmann (2008) study competition between various software providers. In our model, we simplify the structure of the SaaS and on-premises alternatives, using a static model that abstracts away from upgrade cycles and multi-period pricing, in order to elegantly capture software security risk concerns which are the focus of our paper. Specifically, in the following, we will develop an understanding of how negative security externalities and consumer patching behavior interact with SaaS and on-premises offerings to diversify security risk and improve the value associated with software products.

3 Model Description and Equilibrium

A vendor produces software and offers it to a continuum of consumers. The software can be made available in one of two formats: (i) *as a product* to be installed at the consumer's location (on-premises), and (ii) *as a service* which is hosted by the vendor (SaaS). Consumer valuations for the product version are uniformly distributed on $\mathcal{V} = [0, 1]$. However, by opting for the SaaS solution, a consumer loses some flexibility in integration with business systems, ability to control data, and upgrading (Chow et al. 2009). Thus, we assume that if a given consumer has valuation $v \in \mathcal{V}$ for the product version, she has valuation δv for the service version where $0 < \delta < 1$.

We assume that the software is used in a network setting, thereby exposing purchasing consumers to additional risk associated with the software's use. This risk comes in the form of either a directed security attack or an undirected, self-replicating attack (e.g., a worm). We denote the probability that a directed attack occurs on the network with $0 < \pi_d < 1$ (we use d to signify *directed*). Conditional on a directed attack having occurred on the network, we assume that the likelihood any given network location (node) is victimized is proportional to the mass of consumers at that node (Greenemeier and Hoover 2007, Roy 2011). Therefore, the total likelihood of a node that services d consumers being attacked is $\pi_d d$. Similarly, we denote the probability that a patchable security vulnerability arises in the software and that a worm attack on that vulnerability occurs with $0 < \pi_u < 1$ (we use u to signify *undirected*). Given the mechanics of worm spread, if the mass of the unpatched population in the network is u, the unconditional probability that the worm will attack an unpatched user's system is given by $\pi_u u$.

If a user gets struck by either a directed or undirected security attack, then one would expect that she suffers a loss positively correlated with her valuation. That is, consumers with high valuations will incur greater losses than consumers with lower valuations due to higher opportunity costs, higher criticality of data and loss of business. For simplicity, we assume the correlation is of first order, i.e., the loss that a consumer with valuation v incurs if she is hit by the attack is either αv for on-premises or $\alpha \cdot \delta v$ for SaaS, where $\alpha > 0$ is a constant. Undirected attacks typically exploit known vulnerabilities for which a patch is already available, hence each consumer has an opportunity to patch in the face of this security risk.[2] If a consumer chooses to patch the software,

[2] Zero-day attacks on vulnerabilities that do not have a patch available yet can also occur (see, e.g., McBride 2005, IBM 2008), and are central to the debate on software liability because users cannot protect themselves from these

she will incur an expected cost of patching denoted $0 < c_p < 1$, which accounts for the money and effort that a consumer must exert in order to verify, test, and roll-out patched versions of existing systems.

There are three decision periods. In the first period, the vendor determines which versions of its software to release and sets a product price $p > 0$ for a single server license for its on-premises version and a service price $p_s > 0$ for its SaaS version.[3] In the second period, given the price and security risk of each software offering, each consumer makes a decision whether to purchase the software as well as which version to purchase. Finally, in the third decision period, if a patchable security vulnerability has been discovered, each consumer who purchased the on-premises version determines whether or not to patch her own instance. Subsequent to these decision periods, both directed and undirected attacks may realize on the network and consumers incur their respective losses.

Each consumer makes a purchasing decision to buy the on-premises product version, OP, buy the SaaS version, $SaaS$, or not buy either offering, N. Similarly, if a patchable vulnerability arises in the software, each user of the on-premises version makes a decision to either patch, P, or not patch, NP, her own system. We denote the consumer action space by $S = \{OP, SaaS, N\} \times \{P, NP\} - (N, P) - (SaaS, NP)$, the exclusions of (N, P) and $(SaaS, NP)$ stemming from the former clearly being infeasible and the latter reflecting that SaaS implementations need not have patches released to consumers. Given prices, in a consumer market equilibrium, each consumer maximizes her expected utility given the equilibrium strategies of all other consumers. For a strategy profile

risks. In this paper, we focus on patchable vulnerabilities and refer the reader to August and Tunca (2011) which helps build an understanding of how our insights will apply as zero-day attacks become more widespread.

[3] As mentioned in the literature review, we use an essentially static model of product and service offerings in order to focus on security risk issues stemming from consumer usage and patching behavior. We consider each offering to provide value to the consumer for an equivalent fixed amount of time for a fixed price in order to abstract away from many ancillary issues, such as upgrade cycles and dynamic pricing in a dynamic model, while centering in on issues related to security risk diversification. Hence, the price for SaaS (p_s) should be carefully interpreted as the service price for the same period as the server license of the on-premises version.

$\sigma : \mathcal{V} \rightarrow S$, the expected cost faced by the consumer with valuation v is then defined by

$$C(v,\sigma) \triangleq \begin{cases} c_p \quad if \quad \sigma(v) = (OP, P)\,; \\ \pi_u u(\sigma)\alpha v \quad if \quad \sigma(v) = (OP, NP)\,; \\ \pi_d d(\sigma)\alpha\delta v \quad if \quad \sigma(v) = (SaaS, P)\,; \\ 0 \quad if \quad \sigma(v) = (N, NP)\,, \end{cases} \tag{1}$$

where the size of the unpatched user population of the on-premises version is given by

$$u(\sigma) \triangleq \int_{\mathcal{V}} 1_{\{\sigma(v)=(OP,NP)\}}\, dv\,, \tag{2}$$

and the size of the user population of the SaaS version (most vulnerable to a *directed* attack) is given by

$$d(\sigma) \triangleq \int_{\mathcal{V}} 1_{\{\sigma(v)=(SaaS,P)\}}\, dv\,, \tag{3}$$

where $1_{\{\cdot\}}$ is the indicator function. For expositional convenience, we also define the size of the patched population using the on-premises product to be

$$n(\sigma) \triangleq \int_{\mathcal{V}} 1_{\{\sigma(v)=(OP,P)\}}\, dv\,. \tag{4}$$

3.1 Consumer Market Equilibrium

The primary goal of this work is to sharpen our understanding of how introduction of a SaaS version of software affects its corresponding security risk. Therefore, we focus mostly on high security-loss environments where users have stronger incentives to patch their individual installations of on-premises software to avoid economic losses. Before proceeding to the software vendor's pricing problem, we first take prices for the product and service as given and examine the characteristics of the consumer market equilibria that result.

Lemma 1 *Given on-premises and SaaS prices, $p \in (0, 1-c_p)$ and $p_s \in (0, \delta)$, respectively satisfying $p \neq p_s$, and other parameters c_p, π_d, π_u, and δ, there exists $\overline{\alpha} > 0$ such that when $\alpha > \overline{\alpha}$ (i.e., high security-loss environments), a unique equilibrium in the consumer market exists.*[4] *The equilibrium*

[4] Uniqueness is up to a positive measure. When $p = p_s$, it is a special case where consumers are indifferent between

consumer strategy profile σ^* is characterized by thresholds v_d, v_u, $v_p \in [0,1]$ such that for $v \in \mathcal{V}$, it satisfies either

$$\sigma^*(v) = \begin{cases} (OP, P) & if \quad v_p \leq v \leq 1\,; \\ (OP, NP) & if \quad v_u \leq v < v_p\,; \\ (SaaS, P) & if \quad v_d \leq v < v_u\,; \\ (N, NP) & if \quad 0 \leq v < v_d\,, \end{cases} \tag{5}$$

or

$$\sigma^*(v) = \begin{cases} (OP, P) & if \quad v_p \leq v \leq 1\,; \\ (SaaS, P) & if \quad v_d \leq v < v_p\,; \\ (OP, NP) & if \quad v_u \leq v < v_d\,; \\ (N, NP) & if \quad 0 \leq v < v_u\,. \end{cases} \tag{6}$$

In particular, consumer behavior is characterized by three regions in the parameter space:

Region I (No SaaS): If $p_s > \delta c_p$ and $p \leq p_s/\delta - c_p$, then $p < v_u < v_p < 1$ and σ^* is given by either (5) with $v_d = v_u$ or (6) with $v_d = v_p$;

Region II (SaaS for Low-tier): If $p > \max(p_s/\delta - c_p, p_s)$, then $p_s < v_d < v_u < v_p < 1$ and σ^* is given by (5);

Region III (SaaS for Middle-tier): If $p_s < \delta c_p/(1-\delta)$ and $p_s/\delta - c_p \leq p < p_s$, then $p < v_u < v_d < v_p < 1$ and σ^* is given by (6).[5]

Lemma 1 formally establishes that the consumer market exhibits a threshold structure. The patching threshold v_p is the threshold valuation above which consumers prefer to use the on-premises offering and patch when security patches are made available. The threshold v_d marks the valuation above which (up to the next higher threshold) consumers are using SaaS and possibly facing substantial directed risk, and, similarly, v_u marks the valuation above which (again, up to the next higher threshold) consumers are using the on-premises product but not patching, thus exposed to greater undirected security attacks. Notably, in high security-loss environments, the ordering of some market segments can be reversed depending on parameters.

(OP, NP) and $(SaaS, P)$ such that only the sizes of each population need to be maintained in equilibrium; vendor profits will remain unchanged.

[5] For each of the three regions, a complete characterization of the threshold values: v_p, v_u, and v_d, is provided in the Appendix.

In particular, as indicated in Region II and (5) and (6), the middle tier of the consumer market will prefer the on-premises alternative in an unpatched state (OP, NP) when the quality reduction associated with the SaaS offering tends to be larger (i.e., low δ), the SaaS price is high, and the on-premises price is even higher. However, the middle tier will prefer the SaaS alternative otherwise $(SaaS, P)$, unless the price of the on-premises product is sufficiently low that no consumer will opt for SaaS in equilibrium; in this case, there are only two user segments, (OP, P) and (OP, NP), as described in Region I. Lemma 1 also demonstrates that there always exists, in equilibrium, a population of high valuation users who buy and patch the on-premises product in high security-loss environments, i.e., $v_p < 1$.

4 Managing Software Offerings

In the previous section, we established that a consumer market for the SaaS offering may or may not arise depending on prices and other parameters. Furthermore, whether SaaS is targeted to the middle tier or lower tier of the market is similarly dependent. In this section, we formulate the software vendor's pricing problem and study his incentives to price the software offerings to induce these diverse consumer equilibrium outcomes.

Using the segments defined in (2), (3), and (4), the vendor's profit function can be written

$$\Pi(p, p_s) \triangleq p[u(\sigma^*) + n(\sigma^*)] + p_s d(\sigma^*)\,, \tag{7}$$

where the size of each population depends on the equilibrium strategy profile which, in turn, is a function of prices, i.e., $\sigma^* = \sigma^*(\cdot \,|\, p, p_s)$.[6] The vendor's profit maximization problem can then be expressed

$$\begin{aligned} \max_{p, p_s \,\in\, [0,1]} \quad & \Pi(p, p_s) \\ s.t. \quad & (v_d, v_u, v_p) \text{ satisfy } \sigma^*(\cdot \,|\, p, p_s)\,. \end{aligned} \tag{8}$$

In addition to characterizing the optimal prices in (8) and the corresponding equilibrium consumer

[6] In order to focus on the security aspects of the network, we utilize a simplified cost structure with standard assumptions that the software development costs are sunk and the marginal cost of reproduction is sufficiently small to ignore. As for the SaaS alternative, again there is a fixed cost associated with setting up servers and infrastructure and nominal costs associated with servicing the marginal user. Including this type of cost structure will not add significantly to the insights we generate on how security risk can be managed. Hence, we utilize a simplified model here for analysis and discuss how actual SaaS implementation costs would affect our insights in the concluding remarks.

market structures under these prices, we are also interested in examining measures of security risk and social welfare for these outcomes. To facilitate the ensuing discussion, we denote the total security losses with SL and define it as the sum of expected losses from undirected security attacks, directed security attacks, and patching costs under the equilibrium strategy profile σ^*, i.e.,

$$SL \triangleq \int_{\mathcal{V}} 1_{\{\sigma^*(v)\,=\,(OP,NP)\}} \pi_u u(\sigma^*)\alpha v \, dv + \int_{\mathcal{V}} 1_{\{\sigma^*(v)\,=\,(SaaS,P)\}} \pi_d d(\sigma^*)\alpha\delta v \, dv + c_p n(\sigma^*) \,. \tag{9}$$

An appropriate measure for social welfare is given by

$$W \triangleq \int_{\mathcal{V}} \left[1_{\{\sigma^*(v)\in\{(OP,P),\,(OP,NP)\}\}} v + 1_{\{\sigma^*(v)\,=\,(SaaS,P)\}} \delta v \right] dv - SL \,, \tag{10}$$

which is the difference between the aggregate value derived from the software and these losses.

Region I of Lemma 1 states that in high security-loss environments, there exist prices such that the equilibrium consumer market structure is characterized by the absence of a population of consumers using the SaaS software offering. In the following proposition, we demonstrate that in such environments, the vendor strictly prefers to induce equilibrium outcomes in Regions II and III, thus engendering a SaaS user population through his pricing.

Proposition 1 *For high security-loss environments, under optimal pricing there always exists a positive mass of consumers who prefer the SaaS software alternative in equilibrium.*

In high security-loss environments, there can be substantial consumer benefits associated with a reduction in the magnitude of security losses. When the consumer population is induced to separate usage across on-premises and SaaS offerings, these losses are mitigated through diversification. By Lemma 1 and Proposition 1, all three user populations (patched on-premises users, unpatched on-premises users, and SaaS users) are represented in equilibrium. By inducing a population of SaaS users, the vendor has removed a large mass of potentially unpatched hosts from the network, thus reducing the risk faced by remaining unpatched on-premises users. Although the SaaS users are patched and protected from undirected risk, they now face directed risk as a large node on the network. However, because of patching costs, many of these users would not patch as on-premises users, hence pricing the SaaS offerings to induce them to accept some directed risk while completely shielding them from undirected risk helps diversify security risk within the network structure. Proposition 1 establishes that the vendor can efficiently reap the benefits stemming from risk diversification through his pricing; thus, inducing consumers to use SaaS can be profitable.

By Lemma 1 and Proposition 1, it follows that prices will be set such that either $v_p > v_d > v_u$ or $v_p > v_u > v_d$ characterize the consumer market structure in equilibrium. In the former case, the SaaS alternative is targeted at the middle tier of the consumer market (Region III in Lemma 1). Then, by (7), the vendor's profit function is given by

$$\Pi(p, p_s) = p(1 - v_p + v_d - v_u) + p_s(v_p - v_d)\,. \tag{11}$$

We will subsequently refer to the prices that maximize (11) with p^M and p_s^M, constrained such that they induce a middle-tier SaaS consumer market structure. The corresponding profits are denoted by $\Pi^M \triangleq \Pi(p^M, p_s^M)$. Further, security losses satisfy

$$SL = \left[\alpha(\pi_u(v_d - v_u)(v_d^2 - v_u^2) + \delta\pi_d(v_p - v_d)(v_p^2 - v_d^2))\right]/2 + c_p(1 - v_p)\,, \tag{12}$$

and social welfare can be expressed

$$\begin{aligned} W = \big[1 - v_p^2 + v_d^2 - v_u^2 + \delta(v_p^2 - v_d^2) - \alpha(\pi_u(v_d - v_u)(v_d^2 - v_u^2) + \delta\pi_d(v_p - v_d)(v_p^2 - v_d^2))\big]/2 \\ - c_p(1 - v_p)\,. \end{aligned} \tag{13}$$

On the other hand, when the vendor targets its SaaS alternative at the lower tier of the market, inducing $v_p > v_u > v_d$ (Region II of Lemma 1), the vendor's profit function is given by

$$\Pi(p, p_s) = p(1 - v_u) + p_s(v_u - v_d)\,. \tag{14}$$

Similarly, p^L and p_s^L will denote the prices that maximize (14), constrained such that they induce a low-tier SaaS consumer market structure, and the respective profits will be denoted by $\Pi^L \triangleq \Pi(p^L, p_s^L)$. For this structure, the security losses and welfare are given by

$$SL = \left[\alpha(\pi_u(v_p - v_u)(v_p^2 - v_u^2) + \delta\pi_d(v_u - v_d)(v_u^2 - v_d^2))\right]/2 + c_p(1 - v_p)\,, \tag{15}$$

and

$$W = \left[1 - v_u^2 + \delta(v_u^2 - v_d^2) - \alpha(\pi_u(v_p - v_u)(v_p^2 - v_u^2) + \delta\pi_d(v_u - v_d)(v_u^2 - v_d^2))\right]/2 - c_p(1 - v_p)\,, \quad (16)$$

respectively.

Having derived the profit and welfare expressions for each of the feasible equilibrium outcomes established by Proposition 1, we next characterize the conditions under which each outcome is more profitable and how the vendor strategically sets prices to induce them.

Proposition 2 *For high security-loss environments, (i) when patching costs and the SaaS alternative's quality are both high, i.e., $c_p > 1/3$ and $\delta > \frac{2(1-c_p)}{1+c_p}$, a software vendor can maximize profits by setting prices such that the SaaS offering is preferred by the middle tier of the consumer market; (ii) otherwise, the SaaS alternative should be geared for the lower tier of the consumer market.*

First, we discuss part (ii) of Proposition 2. When patching costs are small and security risk is large, consumers have strong incentives to patch when using on-premises software. In this case, the vendor can both charge a high price for its on-premises software and still reduce the security risk faced by unpatched on-premises users because they remain a relatively small population within the network which limits the impact of their security externality. Because of the vendor's pricing power associated with on-premises software, he should keep the price of his SaaS offering lower to prevent cannibalization and serve the lower tier of the consumer market. Part (ii) of Proposition 2 is illustrated in panels (a) and (b) of Figure 1. Because $\delta = 0.80$, the condition $\delta < 2(1 - c_p)/(1 + c_p)$ is satisfied whenever $c_p < 3/7$ as indicated by the area labeled A in the figure. Within this area, the optimal prices are given by p_s^L and p^L satisfying $p_s^L < p^L$ which give rise to a low-tier SaaS structure characterized by $v_p > v_u > v_d$ as illustrated in panel (b) of Figure 1.

For part (ii) of Proposition 2, as patching costs increase (as in the right-hand portion of area A in Figure 1), although there still exist strong incentives to patch due to the high potential security losses, it becomes relatively less affordable to be a patched user of the on-premises software. When patching populations fall, because of the negative security externalities associated with unpatched behavior, overall usage will also decline which reduces vendor profitability. In this case, the vendor must reduce the price of his on-premises product to help maintain a sizeable patching population, as well as to encourage unpatched on-premises users who now face greater risk to remain in the user

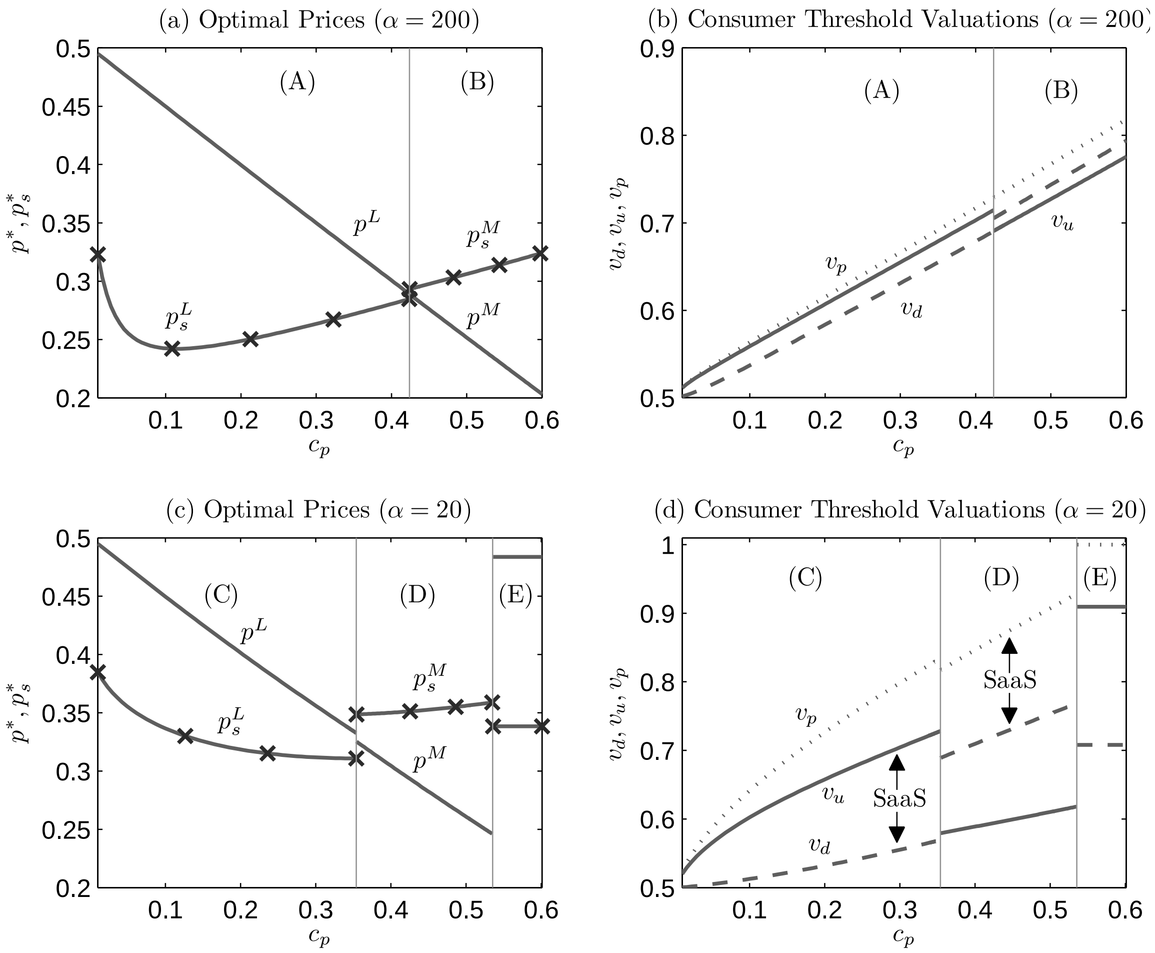

(a) Optimal Prices ($\alpha = 200$)

(b) Consumer Threshold Valuations ($\alpha = 200$)

(c) Optimal Prices ($\alpha = 20$)

(d) Consumer Threshold Valuations ($\alpha = 20$)

Figure 1: How patching costs affect pricing and on-premises versus SaaS usage. The other parameter values for all panels are $\delta = 0.80$, $\pi_u = 0.20$, and $\pi_d = 0.10$.

population. In the left-hand portion of area A in panel (a), the vendor reduces his SaaS price (p_s^L) as patching costs (c_p) increase. This is because when patching costs are small and security risk is high (i.e., $\alpha = 200$), there is a relatively small population of unpatched consumers and the vendor prefers to reduce p_s^L to provide incentives for them to diversify risky usage across on-premises and SaaS varieties. However, as patching costs increase further, it becomes harder for consumers to remain in the market as patched consumers and some shift toward risky usage. In this case, the vendor needs to limit the size of these risky populations (i.e., $u(\sigma^*)$ and $d(\sigma^*)$), by raising p_s^L. When security risk is slightly lower as in panel (c) where $\alpha = 20$, the impact of a slightly larger unpatched, on-premises user population and SaaS user population is not as detrimental, reducing incentives for the vendor to raise p_s^L to restrict risky usage.

Proposition 2 establishes that as patching costs become larger, there exists a point at which the vendor alters his strategy: jumping up his SaaS price from p_s^L to p_s^M, and jumping down his on-premises price from p^L to p^M. This can be seen in panel (a) of Figure 1 as a move from area A (Region II in Lemma 1) to B (Region III in Lemma 1) for a security loss factor of $\alpha = 200$. Similarly, in panel (c) of Figure 1, the same effect is shown as a shift from area C to D for a security loss factor one order smaller. At this point, due to the substantial patching costs, even though he lowers the on-premises product price, the vendor will face a larger unpatched population and reduced usage due to their negative externality. Although his SaaS product may have slightly lower base quality due to less control allotted to users, when accounting for the security externalities, this may not be the case. Specifically, as we cross the boundary in patching costs referred to above, the vendor strategically prices its SaaS product at a higher level, i.e., he targets a smaller higher value population which is accompanied by a smaller directed risk. Because users of the SaaS option are shielded from undirected risk, the vendor's pricing induces an outcome where medium valuation users will prefer the SaaS option over the lower value unpatched on-premises offering that faces more substantial undirected security risk. In panel (b) of Figure 1, the area labeled by B shows how the thresholds induced by his pricing flip to $v_p > v_d > v_u$, leading to a middle-tier SaaS outcome; portion D of panel (d) has a similar nature.

One final point also illustrated in panels (c) and (d) of Figure 1 is another pricing strategy change at the junction between areas D and E. When the magnitude of security losses is not too high (i.e., $\alpha = 20$) and patching costs increase to a larger level, the vendor has incentives to significantly increase his on-premises price to reduce the size of the purchasing on-premises population thus limiting the negative security externality to an extent that these consumers now have much reduced incentives to patch their products. Rather than continuing to cut his on-premises price to ensure a patching population exists to limit undirected security risk, a substantial price increase allows him to serve only the highest valuation market with his on-premises product. Complementing this strategic price increase is a drop in his SaaS price to capture more of the market at the lower end. However, for any level of patching costs, as the security loss factor grows high enough, area E as depicted in Figure 1 disappears due to the large losses users incur when being unpatched; this is the essence of Proposition 2.

Having developed an understanding of the conditions under which the vendor targets his SaaS product to the middle and lower tiers of the consumer market, we next analyze how these outcomes affect social welfare and particularly identify regions where welfare can be increased.

Proposition 3 *For high security-loss environments, when patching costs are within an intermediate range and the SaaS alternative's quality is high, i.e.,* $\underline{c_p} < c_p < 1/3$ *and* $\delta > \underline{\eta}$*, social welfare can be increased if incentives are provided to encourage the software vendor to target his SaaS alternative to the middle tier rather than lower tier of the consumer market. However, for most other levels of patching costs and SaaS quality, the vendor-preferred outcome is also better for welfare. Technically, there exist* $\underline{c_p} > (17 - 4\sqrt{15})/7$ *and* $\underline{\eta} > \frac{2(7-14c_p+3c_p^2)}{(7-c_p)(1+c_p)}$ *such that*

(i) If $\underline{c_p} < c_p < 1/3$ *and* $\delta > \underline{\eta}$*, then* $W\big|_{p^*,p_s^*} < W\big|_{p^M,p_s^M}$*;*

(ii) If $\delta < \frac{2(7-14c_p+3c_p^2)}{(7-c_p)(1+c_p)}$*, or* $c_p > 1/3$ *and* $\delta > \frac{2(1-c_p)}{1+c_p}$*, then* $W\big|_{p^*,p_s^*} = \max\left(W^L, W^M\right)$*,*

where W^L *and* W^M *denote the welfare associated with equilibrium outcomes in Regions II and III, respectively.*[7]

Proposition 3 establishes that there exists an interval of patching costs where the vendor will prefer to induce a low-tier SaaS outcome characterized by $v_p > v_u > v_d$ (Region II of Lemma 1) with his pricing, whereas social welfare would be strictly higher if he priced at p^M and p_s^M to induce the middle-tier SaaS, $v_p > v_d > v_u$ (Region III of Lemma 1), consumer market outcome. The reasoning here is that when the vendor adapts his strategy to targeting the SaaS offering to the middle tier of the consumer market, he effectively increases the size of the patched population by dropping the on-premises price and restricts the size of the SaaS user population by increasing the SaaS price. Combining these effects, the total security losses on the network are smaller which, in aggregate, leads to higher welfare, despite the negative impact of restricted usage. Part (i) of Proposition 3 thus suggests that there may be intervals near the upper bound on patching costs and lower bound on the SaaS quality parameter where providing external incentives to the vendor and/or users to help encourage the socially-preferable outcome. However, part (ii) of the proposition also establishes that, in many cases, an outcome where the SaaS alternative is catered to the lower tier of the consumer market is also consistent with welfare considerations.

As presented in Proposition 1 and further characterized in Propositon 2, in high security-loss environments, the vendor has strong incentives to release a SaaS offering to change the structure

[7] W^L and W^M are defined in the Appendix.

of the network and reduce security risk by splitting the user population; this helps to limit both directed and undirected security attacks which are affected by usage and unpatched population sizes, respectively. Although our focus remains on studying high security-loss environments where the margin for improvement is large, we briefly address an open question of whether versioning (i.e., introducing a SaaS version of an on-premises software product) would make sense for low security-loss environments for which the consumer market structure is characterized in the proof of the following result.

Proposition 4 *For low security-loss environments, a software vendor still prefers to offer both on-premises and SaaS versions of his software.*

In contrast to Proposition 1, one might expect that as the security risk associated with software becomes small that a software vendor would shift toward a strategy where he prices his offerings in such a way that only the higher quality offering is consumed. There is a rich stream of literature on versioning of information goods that explores this topic in detail, examining how the versioning decision relates to cost-to-quality ratios, consumer heterogeneity, positive network effects, competition, asymmetric information, and group tastes (see, e.g., Bhargava and Choudhary 2001, 2008, Johnson and Myatt 2003, Jing 2007, Jones and Mendelson 2011, Niculescu and Wu 2011, Wei and Nault 2011a, 2011b).

Proposition 4 examines versioning from a slightly different perspective. Specifically, in the absence of the negative security externality present in our model, because consumers have uniform valuations and the SaaS offering has a quality reduction factor, the literature cited above informs us that the vendor would optimally offer only the higher quality information good. However, when consumers of both versions are exposed to negative security externalities in the form of directed risk for the SaaS offering and undirected risk for the on-premises offering, Proposition 4 establishes that it is still optimal for the vendor to offer both versions even in low security-loss environments. In fact, as α diminishes, we demonstrate that it is profitable to introduce the SaaS offering. From a practical standpoint, this result provides another alternative explanation for the commonplace existence of multiple versions of software products: as long as the versions have some idiosyncratic risk stemming from their respective user populations, however small, then it is profit-maximizing to set prices such that both versions are consumed in equilibrium.

5 Comparison with Benchmarks

In this section, we examine how a software vendor's decision to release SaaS versions of his traditionally on-premises software product (as detailed in Section 4) affects profitability and the security properties of the network relative to benchmark outcomes where only an on-premises offering is made. Additionally, we study the impact of introducing SaaS on consumer surplus and social welfare. For convenience, we use the subscript "BM" to denote that the measure is under the *benchmark* outcome where the on-premises version is the sole offering.

Propositions 1 and 2 establish that for high security-loss environments, the software vendor will release both alternatives and target the SaaS version to the middle tier when patching costs (c_p) are high and the SaaS version has similar quality (i.e., high δ). In the following proposition, we demonstrate that a joint offering strategy increases both profits and welfare substantially. Further, we characterize how c_p, δ, and the likelihood of a directed attack on the SaaS offering (π_d) affect the extent to which the outcome of the joint offering improves these measures.

Proposition 5 *For high security-loss environments, both vendor profits and social welfare can increase substantially under a joint offering strategy. Both measures are increasing in patching costs and the quality of the SaaS offering, but decreasing in the likelihood of directed attacks. Technically, there exists $\underline{\omega}, \kappa > 0$ such that for all $\alpha > \underline{\omega}$,*

$$\left| \frac{\Pi^* - \Pi_{BM}}{\Pi_{BM}} - \frac{c_p \delta}{\pi_d \alpha (1 - c_p)^2} \right| < \frac{\kappa}{\alpha^2} \tag{17}$$

and

$$\left| \frac{W^* - W_{BM}}{W_{BM}} - \frac{2 c_p \delta}{3 \pi_d \alpha (1 - c_p)^2} \right| < \frac{\kappa}{\alpha^2} \tag{18}$$

are satisfied.

Proposition 5 establishes that the introduction of a SaaS offering can result in substantial percentage increases in profits and social welfare. Examining the inequalities in (17) and (18), it is straightforward to see that both normalized measures decrease in π_d but increase in c_p and δ. A decrease in π_d corresponds to reduced directed security risk for consumers who use the SaaS alternative in equilibrium. In a similar vein, an increase in δ also reflects a higher quality SaaS offering which is beneficial to both vendor profitability and welfare. On the other hand, for c_p, the potential improvement associated with a SaaS release stems from consumers' patching behavior of the on-premises solution. In particular, as patching costs increase, consumers find it incentive compatible to bear more undirected security risk rather than incurring these patching costs. Given the negative externalities unpatched users can inflict on the network, under these circumstances, introducing the SaaS alternative can have an even stronger effect by inducing consumers to split usage across alternatives and diversify security risk.

For low security-loss environments, we characterize the relative benefit of introducing SaaS in the following proposition.

Proposition 6 *For low security-loss environments, expansion to a SaaS offering will provide a limited increase in vendor profits and social welfare. The associated benefits are increasing in the quality of the SaaS offering and the likelihood of undirected attacks. Technically, there exists $\overline{\omega}, \eta > 0$ such that for all $\alpha < \overline{\omega}$,*

$$\left| \frac{\Pi^* - \Pi_{BM}}{\Pi_{BM}} - \frac{\delta \pi_u^2 \alpha^2}{16(1-\delta)} \right| < \eta \alpha^3 \tag{19}$$

and

$$\left| \frac{W^* - W_{BM}}{W_{BM}} - \frac{5\delta \pi_u^2 \alpha^2}{48(1-\delta)} \right| < \eta \alpha^3 \tag{20}$$

are satisfied.

Although Proposition 4 demonstrates that a vendor should optimally release both on-premises and SaaS versions of his product in low security-risk environments, Proposition 6 suggests that the benefits stemming from diversification are much more limited in these environments. In contrast to Proposition 5, by comparing (17) and (19), the percentage increase in profits associated with releasing the SaaS alternative is an order of magnitude smaller and may not justify the additional costs of managing two versions of the software.

The findings in Propositions 5 and 6 are illustrated in Figure 2. We depict three curves plotting the measure, $\frac{\Pi^* - \Pi_{BM}}{\Pi_{BM}}$, computed numerically under parameter sets A: $c_p = 0.30, \pi_u = 0.23$, B: $c_p' = 0.50, \pi_u = 0.23$, and C: $c_p = 0.30, \pi_u' = 0.55$, respectively. As can be seen, near $\alpha = 30$, the percentage improvement in profits ranges from approximately 10-30%; however, near $\alpha = 1/30$, the percentage improvement is negligible. This is the essence of the two propositions - that diversifica-

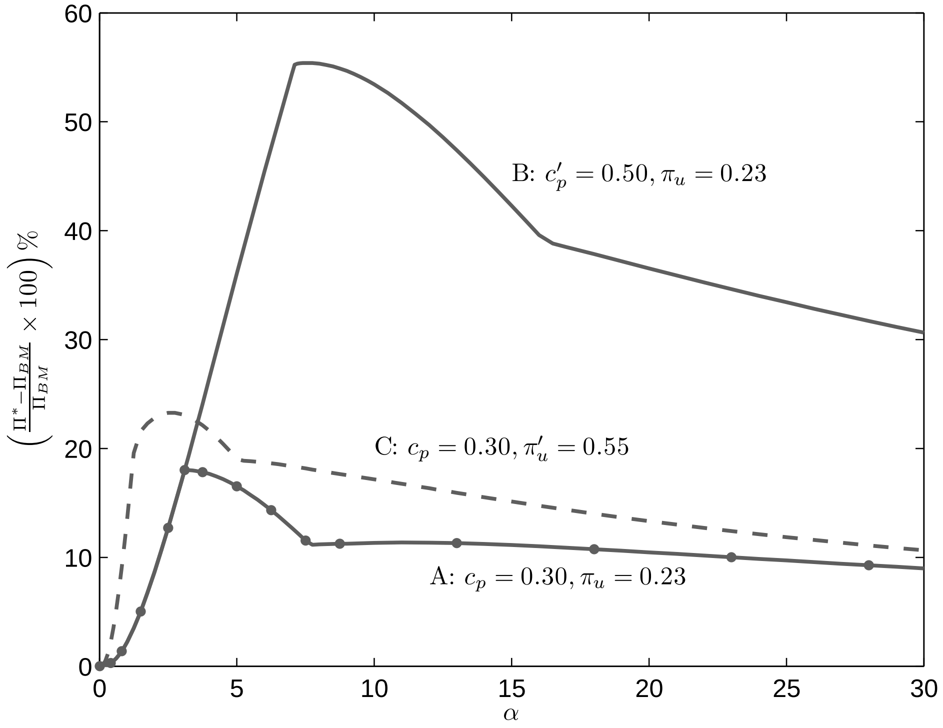

Figure 2: Percentage increase in vendor profitability when a SaaS version is offered in addition to an on-premises version of software. This percentage is plotted over a wide range of security loss environments. Values for patching costs and undirected attack probabilities are listed on the plot. The other parameter values are $\delta = 0.80$ and $\pi_d = 0.10$.

tion of security risk by offering SaaS has much greater potential for moderate to high security-loss environments. Comparing curve B to A, one can see that an increase in patching costs from $c_p = 0.30$ to $c_p' = 0.50$ can change the potential profit improvement of a SaaS release strategy up to over 40% because of the poor patching behavior induced on the network at this higher cost level.[8] This characteristic is consistent with the results presented in Proposition 5, in particular from (17). Similarly, for lower α, Proposition 6 and specifically (19) suggests that an increase in the likelihood of an undirected attack (π_u) will also increase the potential benefit of a diversification strategy. In Figure 2, we can see this effect by comparing curve C to A at the lower range of α.

Next, focusing our attention on high security-loss environments which have the greatest potential for improvement, we examine how introducing a SaaS alternative affects the security properties of the network as well as consumer surplus. By (9), we can define the average per-user security losses

[8] When the security loss factor (α) is within a lower range, curves A and B coincide because the equilibrium consumer market structure under optimal pricing dictates that users of the on-premises product are not patching, preferring to bear the low security risk.

as

$$\hat{SL} \triangleq \frac{SL}{u(\sigma^*) + d(\sigma^*) + n(\sigma^*)}\,, \tag{21}$$

which simplifies to either $SL/(1-v_d)$ or $SL/(1-v_u)$ in Regions II and III of Lemma 1, respectively.

Proposition 7 *When a SaaS offering is introduced in high security-loss environments, the average security losses per user decrease under high patching costs, i.e. $c_p \geq \delta/(4-\delta)$, but actually increase otherwise.*

Proposition 7 brings forth an important insight: a vendor's diversification of software usage by offering both on-premises and SaaS varieties can actually increase per-user security losses. One would expect that introducing a SaaS alternative would split the undirected risk being faced in the benchmark case into two smaller risks (undirected and directed) as a portion of the consumers adopt the SaaS alternative instead. However, Proposition 7 establishes that a software vendor may influence usage and patching behavior through pricing in such a way that the average security losses per user is higher in the joint offering.

In high security-loss environments, when SaaS is introduced, some consumers who would have elected to buy the on-premises product and remain unpatched, i.e., (OP, NP), in the benchmark case now have incentives to switch to SaaS usage, i.e., $(SaaS, P)$. Because this reduces the size of the unpatched population, consumers who were buying the on-premises product and patching, i.e., (OP, P), are now no longer facing as large a negative externality. Therefore, they have overall reduced incentives to patch, and some of these consumers will now elect to remain unpatched instead. Also, because introduction of SaaS splits risk into undirected and directed types, some consumers who had opted out in the benchmark case will now become users. Thus, in comparison to the benchmark case, when both on-premises and SaaS versions are offered, overall usage increases while overall patching decreases.

Proposition 7 establishes that when patching costs are low, the aforementioned cumulative effect of increased usage and decreased patching associated with the introduction of SaaS results in higher average per-user security losses. The reason is that when patching costs are small, the consumer market structure is already characterized by a large patching population in the benchmark case. The population of unpatched on-premises users is in contrast relatively small. Hence, when SaaS is introduced, although patching slightly decreases, the proportional increase in either unpatched or

SaaS usage is substantial. A relatively large increase in these two types of usage which are exposed to undirected and directed security risk, respectively, can lead to higher average security losses because of the accompanying negative externalities. On the other hand, when patching costs are large, the benchmark case is characterized by a small patching population and large unpatched population. In this case, aggregate SaaS and unpatched on-premises usage still increases while patching decreases. However, the reduction in patching behavior has a relatively minor negative effect on an already substantial unpatched population. In contrast to the case above, the diversification benefits of splitting security risk into undirected and directed types now outweigh the minor increase in the externality. Thus, for large patching costs, per-user security losses decrease when SaaS is made available.

Because, surprisingly, average per-user security losses can increase when a software vendor introduces a SaaS version of his on-premises product, from a consumer perspective such a release may not necessarily be beneficial.

Proposition 8 *Despite the substantial increase in welfare stemming from a SaaS release under high security-loss environments, when patching costs are low, i.e.,* $c_p \leq 1/3$*, and the SaaS offering quality parameter satisfies* $\delta > 2 - \frac{64c_p^2\pi_d(1-c_p(4-c_p))}{\pi_u(1+c_p)^4}$*, consumer surplus decreases in equilibrium.*

Proposition 8 suggests that for software that has relatively lower patching costs (such as simple client application software as opposed to enterprise server software), a vendor will release a SaaS version not for the benefits of reduced security risk but rather to expand his market at the lower end and price discriminate. The net effect of his joint offerings on consumer surplus is negative, which is partly driven by the increase in security losses formalized in Proposition 7. Thus, it is useful to characterize how a risk diversification strategy compares in terms of security losses and social welfare to other basic means of improving software security. In the next proposition, we study their comparative performance.

Proposition 9 *From a social planner's perspective, campaigns to reduce undirected attack likelihood and patching costs can be more effective than diversifying security risk through versioning. Specifically, for sufficiently low* α*, a moderate reduction in* π_u *can reduce average per-user security losses and increase welfare to a greater extent than the addition of SaaS to the vendor's offerings. For sufficiently high* α*, a moderate reduction in* c_p *similarly outperforms versioning.*

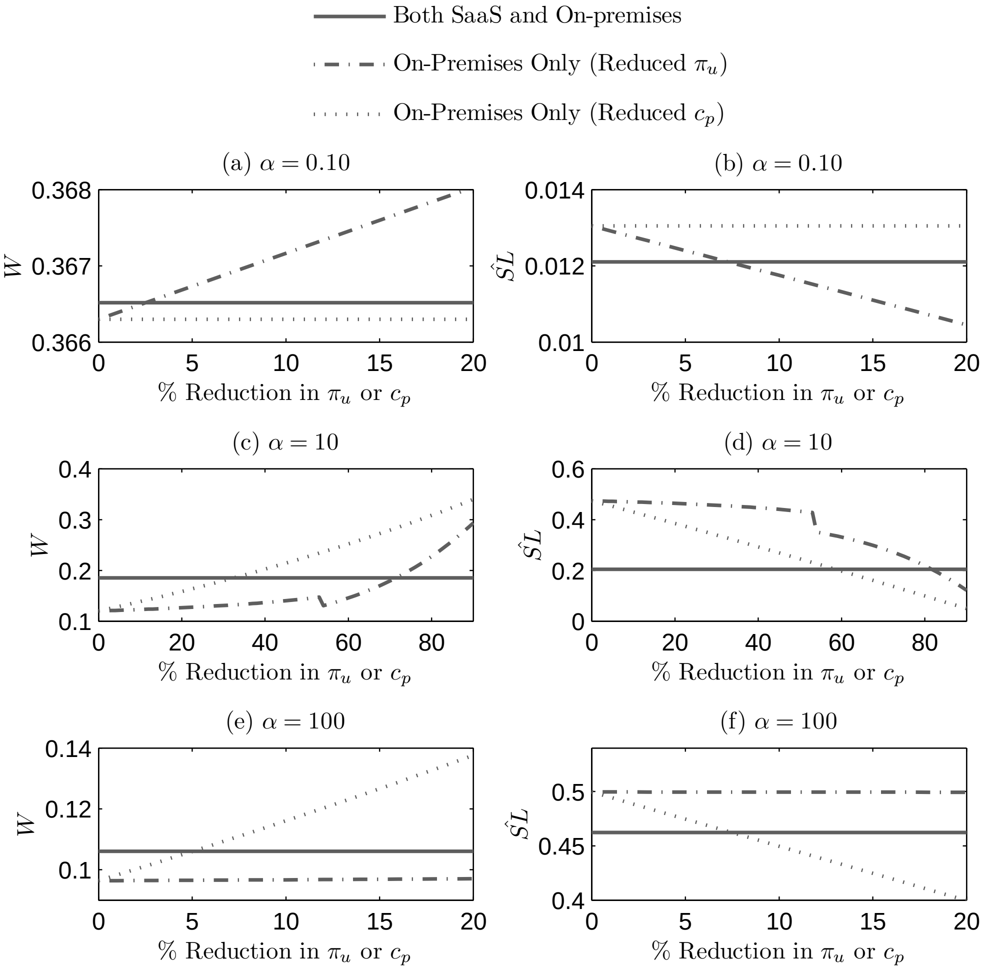

Figure 3: The effect of releasing a SaaS alternative on welfare and per-user security losses compared to reductions in the likelihood of undirected attacks and patching costs. For all panels, $\pi_u = 0.35$ and $c_p = 0.50$, and the x-axis reflects percentage reductions in these nominal values for the dotted and dash-dotted curves. The security loss factor (α) ranges over three levels (0.10, 10, and 100), and the other parameter values are $\delta = 0.80$ and $\pi_d = 0.10$.

Proposition 9's results are illustrated in Figure 3 which compares welfare and per-user security losses across three categories: diversification of risk through SaaS versioning, reduction in security attack likelihood, and reduction in patching costs. Panels (a), (c), and (e) depict the impact of these three alternative software security improvements on social welfare, while panels (b), (d), and (f) demonstrate it for per-user security losses. For low security-loss environments, panel (a) illustrates that a small reduction in π_u (around 3%) will already increase welfare to a greater extent than a diversification strategy using SaaS. Similarly, as panel (b) shows, for a slightly greater reduction in π_u (around 7%), per-user security losses can also be lower than the SaaS case. These numerical results are consistent with what is established in Proposition 6, namely that the benefits of security risk diversification is limited in low security-loss environments. On the other hand, for rather high security-loss environments, panels (e) and (f) reflect a similar finding that going directly after a small reduction in patching costs (c_p) helps the most. In this case, Proposition 5 establishes that the network can benefit greatly from SaaS release and diversification of security risk. However, as the potential for losses becomes critically high, the user population is characterized by a large number of patched, on-premises users. In such instances, reducing patching costs can increase usage substantially because more consumers can suddenly afford to use and shield themselves from undirected risk. This behavior is consistent with the results in Proposition 9.

What Figure 3 also depicts is that for security-loss environments that are within a moderate to slightly high range, although reductions in patching costs can still improve welfare and reduce security losses, the magnitude of the reductions would have to be quite large for this strategy to outperform a risk diversification one. In panels (c) and (d), where $\alpha = 10$, patching costs would need to be reduced by over 35% and 55%, to perform better than SaaS versioning in terms of social welfare and security losses, respectively. Therefore, from a social planner's perspective, versioning for risk diversification is most effective for these moderate to slightly high security-loss environments where it may be difficult to find ways to extensively reduce patching costs.

6 Concluding Remarks

In this paper, we explored the impact of SaaS offerings on software security risks. In particular, we focused on examining how the inclusion of SaaS affects consumer preferences for both SaaS and on-premises varieties and how a shift in the usage of each alternative, in turn, affects patching incentives for on-premises users. Based upon our characterization of equilibrium consumer behavior that accounts for negative security externalities stemming from unpatched on-premises usage and aggregate SaaS usage, we analyzed how software vendors can use SaaS offerings to create benefits from diversified security risk.

We showed that in high security-loss environments, a software vendor has strong incentives to release a SaaS version of his on-premises software product and set prices such that a significant portion of consumers opt for SaaS in equilibrium. When both patching costs and the relative quality of the SaaS offering are high, the vendor should gear the SaaS offering to the middle tier of the market; otherwise, it should be catered to the lower tier. However, social welfare can be higher under middle-tier SaaS outcomes so social planners may want to provide additional incentives to encourage them. We also established that for low security-loss environments, offering SaaS is still beneficial to the software vendor. This result contributes to the information goods versioning literature by establishing that it is still optimal to version as long as separate software versions have idiosyncratic risk stemming from their respective user populations (here, in the form of a negative security externality), even as this risk becomes small.

We demonstrated that the potential improvement to profits and social welfare associated with a SaaS release and its corresponding security risk diversification are substantial in high security-loss environments, but limited in low security-loss ones. When patching costs are high, the average per-user security losses decrease as a result of this diversification. In contrast, when patching costs are low, we show that average per-user security losses can actually increase, and further, if the quality of the SaaS offering is relatively high, consumer surplus can decrease as a result in equilibrium. This suggests that although SaaS can help diversify security risk, under low patching costs, the vendor has profit incentives to increase user populations to an extent where security externalities also grow and leave consumers worse off. Finally, we established that a security risk diversification strategy through SaaS has the best promise for moderate to slightly high security-loss environments where it would take a significantly large reduction in either patching costs or security attack likelihood to attain a comparable outcome with regard to social welfare and security loss measures. For low security-loss environments, a small reduction in the attack likelihood can yield the greatest benefit.

In recent years, companies have invested millions of dollars to provide cloud computing services. In order to implement SaaS alternatives in addition to their traditional on-premises offerings, software vendors would necessarily need to incur additional costs in practice. These costs would have a fairly large fixed-variable cost ratio. In that sense, all of our model implications (e.g., when SaaS is most effective, how it should be targeted to consumers, etc.) remain the same provided that the increased profitability stemming from SaaS covers the fixed costs of providing these services. Particularly, our results indicate that the actual implementation costs are easiest to justify in moderately high security-loss environments where SaaS usage populations are substantial.

The benefits of cloud computing and, particularly, SaaS applications are driving businesses and governments to migrate many internally supported systems and software to the cloud. However, such a paradigm shift can have major consequences on security risks as consumers make choices on software deployment and protection. We hope that the model and insights presented in this paper provide a stepping stone toward a broader understanding of how security risk can be managed as cloud computing matures and becomes an integral part of IT strategies.

Acknowledgments

The authors thank Sridhar Narasimhan, D. J. Wu, Subodha Kumar, Ali Tafti for helpful comments and discussions. This material is based upon work supported by the National Science Foundation under Grant No. CNS-0954234.

References

Anderson, R. and T. Moore (2006). The economics of information security. *Science 314*, 610–613.

Anderson, R. J. (2001). Why information security is hard – an economic perspective. In *Proc. of the 17th Annual Computer Security Applications Conf.*, pp. 358–365. IEEE Computer Soc.

Arora, A., R. Telang, and H. Xu (2008). Optimal policy for software vulnerability disclosure. *Management Science 54*(4), 642–656.

Aspan, M. and C. Baldwin (2011, Apr). Sony breach could cost card lenders $300 million. *Reuters.*

August, T. and T. I. Tunca (2006). Network software security and user incentives. *Management Science 52*(11), 1703–1720.

August, T. and T. I. Tunca (2008). Let the pirates patch? An economic analysis of software security patch restrictions. *Information Systems Research 19*(1), 48–70.

August, T. and 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.

Bain, C., D. B. Faatz, A. Fayad, and D. Williams (2002). Diversity as a defense strategy in information systems. Does evidence from previous events support such an approach? In *Proceedings of the IFIP TC11/WG11.5 Fourth Working Conference on Integrity, Internal Control and Security in Information Systems: Connecting Governance and Technology*, Deventer, The Netherlands, The Netherlands, pp. 77–94. Kluwer, B.V.

Bhargava, H. K. and V. Choudhary (2001). Information goods and vertical differentiation. *Journal of Management Information Systems 18*(2), 89–106.

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

Cavusoglu, H., H. Cavusoglu, and S. Raghunathan (2007). Efficiency of vulnerability disclosure mechanisms to disseminate vulnerability knowledge. *IEEE Transactions on Software Engineering 33*(3), 171–185.

Cavusoglu, H., H. Cavusoglu, and J. Zhang (2008). Security patch management: Share the burden or share the damage? *Management Science 54*(4), 657–670.

Chen, P., G. Kataria, and R. Krishnan (2011). Correlated failures, diversification, and information security risk management. *MIS Quarterly 35*(2), 397–422.

Choi, J. P., C. Fershtman, and N. Gandal (2010). Network security: Vulnerabilities and disclosure policy. *Journal of Industrial Economics 58*(4), 868–894.

Choudhary, V. (2007). Comparison of software quality under perpetual licensing and software as a service. *Journal of Management Information Systems 24*(2), 141–165.

Chow, R., P. Golle, M. Jakobsson, E. Shi, J. Staddon, R. Masuoka, and J. Molina (2009). Controlling data in the cloud: outsourcing computation without outsourcing control. In *Proceedings of the 2009 ACM workshop on Cloud computing security*, CCSW '09, New York, NY, USA, pp. 85–90. ACM.

Claburn, T. (2009, Sep). Government embraces cloud computing, launches app store. *InformationWeek*.

Dey, D., A. Lahiri, and G. Zhang (2012). Hacker behavior, network effects, and the security software market. Forthcoming in *Journal of Management Information Systems*.

Efrati, A. and S. Gorman (2011, June). Google mail hack blamed on China. *The Wall Street Journal*.

El Akkad, O. (2011, Jun). Microsoft wants you to rent an Office in the cloud. *The Globe and Mail*.

Espiner, T. (2007, Nov). Salesforce tight-lipped after phishing attack. *ZDNet*.

Gordon, L. A. and M. P. Loeb (2002, November). The economics of information security investment. *ACM Trans. Inf. Syst. Secur. 5*, 438–457.

Gorman, S. and J. E. Vascellaro (2010, Feb). Google attack linked to Asian hackers. *The Wall Street Journal*.

Greenemeier, L. and J. N. Hoover (2007, Feb). How does the hacker economy work? *InformationWeek*.

Huang, K. and A. Sundararajan (2005, Nov). Pricing models for on-demand computing. National University of Singapore, New York University.

IBM (2008). IBM internet security systems X-Force 2008 mid-year trend statistics. *IBM Global Technology Services*.

Jing, B. (2007). Network externalities and market segmentation in a monopoly. *Economics Letters 95*, 7–13.

Johnson, J. P. and D. P. Myatt (2003, Jun.). Multiproduct quality competition: Fighting brands and product line pruning. *The American Economic Review 93*(3), 748–774.

Johnson, M. E. (2008). *Managing Information Risk and the Economics of Security* (1st ed.). Springer Publishing Company, Incorporated.

Jones, R. and H. Mendelson (2011). Information goods vs. industrial goods: Cost structure and competition. *Management Science 57*(1), 164–176.

Keizer, G. (2004, May). Sasser worm impacted businesses around the world. *TechWeb News*.

Keizer, G. (2008, Dec). Windows users indifferent to Microsoft patch alarm, says researcher. *Computerworld*.

Kim, B. C., P. Chen, and T. Mukhopadhyay (2010). An economic analysis of the software market with a risk-sharing contract. *International Journal of Electronic Commerce 14*(2), 7–39.

Kim, B. C., P. Chen, and T. Mukhopadhyay (2011). The effect of liability and patch release on software security: The monopoly case. *Production and Operations Management 20*(4), 603–617.

Kundra, V. (2011, Aug). Tight budget? Look to the ‘cloud’. *The New York Times*.

Lahiri, A. (2011). Revisiting the incentive to tolerate illegal distribution of software products. *Hawaii International Conference on System Sciences*, 1–9.

Lemos, R. (2003, Feb). ‘Slammer’ attacks may become way of life for net. *CNET News.com*.

Lemos, R. (2004, Apr). MSBlast epidemic far larger than believed. *CNET News.com*.

Ma, D. and A. Seidmann (2008, Aug.). Pricing on-demand software competitively in a dynamic market. Singapore Management University, University of Rochester.

Markoff, J. (2009, Aug). Defying experts, rogue computer code still lurks. *The New York Times*.

McBride, S. (2005, Feb). Zero day attack imminent. *Computerworld*.

Moore, D., C. Shannon, and J. Brown (2002). Code-Red: a case study on the spread and victims of an internet worm. *Proceedings of the Second ACM SIGCOMM Workshop on Internet Measurement*, 273–284.

Niculescu, M. F. and D. Wu (2011, May). When should software firms commercialize new products via freemium business models? Working Paper, Georgia Institute of Technology.

O’Neill, S. (2011, Sept). Survey: Value of the cloud, telecommuting overstated. *CIO*.

Osawa, J. (2011, May). As Sony counts hacking costs, analysts see billion-dollar repair bill. *The Wall Street Journal*.

Png, I. P. and Q. Wang (2009). Information security: Facilitating user precautions vis-à-vis enforcement against attackers. *Journal of Management Information Systems 26*(2), 97–121.

Rahman, M. S., K. Kannan, and M. Tawarmalani (2007, Mar). The countervailing incentive of restricted patch distribution: Economic and policy implications. University of Calgary, Purdue University.

Ransbotham, S. and S. Mitra (2009). Choice and chance: A conceptual model of paths to information security compromise. *Information Systems Research 20*(1), 121–139.

Ransbotham, S., S. Mitra, and J. Ramsey (2010). Are markets for vulnerabilities effective? Forthcoming in *MIS Quarterly*.

Roy, D. (2011, Aug). Data on sale. *CIO*.

Sherr, I. and N. Wingfield (2011, May). Play by play: Sony’s struggles on breach. *The Wall Street Journal*.

Wei, X. and B. R. Nault (2011a, Jan.). Monopoly versioning of information goods when consumers have group tastes. Fudan University, University of Calgary.

Wei, X. and B. R. Nault (2011b, Feb.). Vertically differentiated information goods: Monopoly power through versioning. Fudan University, University of Calgary.

Zetter, K. (2010, Jan). Hack of Google, Adobe conducted through zero-day IE flaw. *WIRED.*

Zhang, J. J. and A. Seidmann (2010). Perpetual versus subscription licensing under quality uncertainty and network externality effects. *Journal of Management Information Systems 27*(1), 39–68.

Online Supplement for

"Cloud Implications on Software Network Structure and Security Risks"

Proof of Lemma 1: First, $\sigma(v) = (OP, P)$ if and only if

$$v \geq \max\left(\frac{c_p}{\pi_u \alpha u}, \frac{p + c_p - p_s}{1 - \delta(1 - \pi_d \alpha d)}, p + c_p\right) . \tag{A.1}$$

By (A.1), in equilibrium, if a consumer with valuation v_0 buys and patches the on-premises alternative, then every consumer with valuation $v > v_0$ will also do so. Hence, there exists a $v_p \in (0, 1]$ such that for all $v \in \mathcal{V}$, $\sigma^*(v) = (OP, P)$ if and only if $v \geq v_p$. Similarly, $\sigma(v) \in \{(OP, P), (OP, NP), (SaaS, P)\}$, which is to say the consumer purchases one of the alternatives, if and only if

$$v \geq \min\left(\frac{p}{1 - \pi_u \alpha u}, \frac{p_s}{\delta(1 - \pi_d \alpha d)}, p + c_p\right) . \tag{A.2}$$

Let $0 < v_1 \leq 1$ and $\sigma^*(v_1) \in \{(OP, P), (OP, NP), (SaaS, P)\}$, then by (A.2), for all $v > v_1$, $\sigma^*(v) \in \{(OP, P), (OP, NP), (SaaS, P)\}$, and hence there exists a $\underline{v} \in (0, 1]$, such that a consumer with valuation $v \in \mathcal{V}$ will purchase if and only if $v \geq \underline{v}$.

By (A.1) and (A.2), $\underline{v} \leq v_p$. A purchasing consumer with valuation v will prefer (OP, NP) over $(SaaS, P)$ if and only if $v[(1 - \pi_u \alpha u) - d(1 - \pi_d \alpha d)] > p - p_s$. If this inequality is satisfied for some $\underline{v} < v_2 \leq 1$, then it is satisfied for all $v > v_2$. It follows that there exists $v_u \in [\underline{v}, v_p]$ such that $\sigma(v) = (OP, NP)$ for all $v \in [v_u, v_p]$, and $\sigma(v) = (SaaS, P)$ for all $v \in [v_d, v_u]$ where $v_d = \underline{v}$. The other possible ordering, i.e. $v_u \leq v_d$, can similarly be shown which establishes the threshold structure in (5) and (6).

Characterizing the bounds of Regions I–III is straightforward analysis closely following that found in August and Tunca (2006, 2008); here we omit the case-by-case details for brevity and provide the characterization of the actual threshold valuations. For Region I, v_u satisfies

$$v_u = \sup\{v_u \mid \pi_u \alpha v_u^2 (v_u - c_p - p) = -(v_u - p)^2\} , \tag{A.3}$$

and

$$v_p = \frac{c_p v_u}{v_u - p}\,. \tag{A.4}$$

For Region II, v_d, v_p, and v_u satisfy

$$v_d = \frac{-\delta(1 - v_u\alpha\pi_d) + \sqrt{4p_s\alpha\delta\pi_d + \delta^2(1 - v_u\alpha\pi_d)^2}}{2\alpha\delta\pi_d}\,, \tag{A.5}$$

$$v_p = \frac{v_u\alpha\pi_u + \sqrt{\alpha\pi_u(4c_p + v_u^2\alpha\pi_u)}}{2\alpha\pi_u}\,, \tag{A.6}$$

and

$$\begin{aligned} v_u = \sup\Big\{v_u \,\Big|\, & 2(p - p_s) + v_u\sqrt{\delta(4p_s\alpha\pi_d + \delta(1 - v_u\alpha\pi_d)^2)} + \\ & v_u\sqrt{\alpha\pi_u(4c_p + v_u^2\alpha\pi_u)} + v_u(\delta - 2 - v_u\alpha(\delta\pi_d + \pi_u)) = 0\Big\}\,, \end{aligned} \tag{A.7}$$

respectively. Finally, for Region III, v_u, v_p, and v_d satisfy

$$v_u = \frac{-1 + v_d\alpha\pi_u + \sqrt{4p\alpha\pi_u + (1 - v_d\alpha\pi_u)^2}}{2\alpha\pi_u}\,, \tag{A.8}$$

$$v_p = \frac{-1 + \delta + v_d\alpha\delta\pi_d + \sqrt{4(p + c_p - ps)\alpha\delta\pi_d + (1 - \delta - v_d\alpha\delta\pi_d)^2}}{2\alpha\delta\pi_d}\,, \tag{A.9}$$

and

$$\begin{aligned} v_d = \sup\Big\{v_d \,\Big|\, & 2(p_s - p) + v_d\sqrt{4(c_p + p - p_s)\alpha\delta\pi_d + (-1 + \delta + v_d\alpha\delta\pi_d)^2} + \\ & v_d\sqrt{4p\alpha\pi_u + (-1 + v_d\alpha\pi_u)^2} - v_d(\delta + v_d\alpha\delta\pi_d + v_d\alpha\pi_u) = 0\Big\}\,, \end{aligned} \tag{A.10}$$

respectively. ■

Proof of Proposition 1: The analysis for this case can be found in the proof of Proposition 5 which examines the benchmark case - an equivalent outcome to Region I of Lemma 1. Comparing the optimal profit expression in the benchmark case to the measures derived for Regions II and III in the proof of Proposition 2, the result follows. ■

Proof of Proposition 2: Technically, we will prove that there exists $\overline{\alpha} > 0$ such that when $\alpha > \overline{\alpha}$, p^* and p_s^* are set so that

(i) If $c_p > 1/3$ and $\delta > \frac{2(1-c_p)}{1+c_p}$, then $\sigma^*(v)$ is characterized by $1 > v_p > v_d > v_u > 0$ and given in (6);

(ii) Otherwise, $\sigma^*(v)$ is characterized by $1 > v_p > v_u > v_d > 0$ and given in (5).

By part (i) of Lemma 1 and Proposition 1, for sufficiently high α, either $v_p > v_u > v_d$ or $v_p > v_d > v_u$ is satisfied under optimal pricing in equilibrium. Suppose $v_p > v_d > v_u$. By (A.10), we obtain

$$v_d = p + c_p + \frac{pc_p\delta\pi_d - \pi_u(p + c_p - p_s)(\delta(p + c_p) - p_s)}{\alpha\delta\pi_d\pi_u(p + c_p)^2} + O\left(\frac{1}{\alpha^2}\right). \tag{A.11}$$

Substituting (A.11) into (A.8) and (A.9), we obtain

$$v_u = p + c_p - \frac{c_p^2\delta\pi_d + \pi_u(p + c_p - p_s)(\delta(p + c_p) - p_s)}{\alpha\delta\pi_d\pi_u(p + c_p)^2} + O\left(\frac{1}{\alpha^2}\right), \tag{A.12}$$

and

$$v_p = p + c_p + \frac{pc_p\delta\pi_d + \pi_u p_s(\delta(p + c_p) - p_s)}{\alpha\delta\pi_d\pi_u(p + c_p)^2} + O\left(\frac{1}{\alpha^2}\right). \tag{A.13}$$

Substituting (A.11), (A.12), and (A.13) into (11) which applies in this case, we obtain

$$\Pi(p, p_s) = p(1 - p - c_p) + \frac{pc_p^2\delta\pi_d + \pi_u c_p p_s(\delta(p + c_p) - p_s)}{\alpha\delta\pi_d\pi_u(p + c_p)^2} + O\left(\frac{1}{\alpha^2}\right). \tag{A.14}$$

By (8) and (A.14), the interior maximizing prices satisfy

$$p^M = \frac{1 - c_p}{2} + \frac{2c_p^2(3c_p - 1)}{\pi_u\alpha(1 + c_p)^3} + O\left(\frac{1}{\alpha^2}\right), \tag{A.15}$$

and

$$p_s^M = \frac{\delta(1 + c_p)}{4} + \frac{A_1}{16c_p(1 + c_p)^3\pi_d\pi_u\alpha} + O\left(\frac{1}{\alpha^2}\right), \tag{A.16}$$

where $A_1 = 8c_p\pi_d(3 - c_p^2(3 + 2\delta) + c_p(2\delta - 5) + c_p^3(5 + 4\delta)) + \delta\pi_u(1 + c_p)^3(1 + c_p(2\delta - 3))$. By substituting (A.15) and (A.16) into (A.14), the vendor's profits are given by

$$\Pi^M = \Pi(p^M, p_s^M) = \frac{(1 - c_p)^2}{4} + \frac{c_p}{4\alpha}\left(\frac{\delta}{\pi_d} + \frac{8c_p(1 - c_p)}{\pi_u(1 + c_p)^2}\right) + \frac{64c_p^3\pi_d^2 A_2 + (1 + c_p)^3(\delta - 2)\pi_u A_3}{16(1 + c_p)^6\pi_d^2\pi_u^2\alpha^2} + O\left(\frac{1}{\alpha^3}\right). \tag{A.17}$$

where $A_2 = -4 + 5c_p + 5c_p^3 - 2c_p^2$ and $A_3 = 8c_p\pi_d(3 - c_p)(1 - c_p) + \delta\pi_u(1 + c_p)^3$.

On the other hand, suppose $v_p > v_u > v_d$. Following similar analysis to the above and using (14), the interior maximizing prices satisfy

$$p^L = \frac{1-c_p}{2} + \frac{2c_p^2(3c_p-1)}{\pi_u\alpha(1+c_p)^3} + O\left(\frac{1}{\alpha^2}\right), \tag{A.18}$$

and

$$p_s^L = \frac{\delta(1+c_p)}{4} + \left(\frac{\delta(1+c_p(2\delta-3))}{16c_p\pi_d\alpha} - \frac{c_p^2\delta(3-c_p)}{(1+c_p)^3\pi_u\alpha}\right) + O\left(\frac{1}{\alpha^2}\right). \tag{A.19}$$

Substituting (A.18) and (A.19) into (14), we obtain

$$\Pi^L = \Pi(p^L, p_s^L) = \frac{(1-c_p)^2}{4} + \frac{c_p}{4\alpha}\left(\frac{\delta}{\pi_d} + \frac{8c_p(1-c_p)}{\pi_u(1+c_p)^2}\right) + \frac{A_4}{\pi_u\alpha^2(1+c_p)^3} + O\left(\frac{1}{\alpha^3}\right). \tag{A.20}$$

where

$$A_4 = \left(\frac{\delta(\pi_u(\delta-2)(1+c_p)^3 + 16c_p^2\pi_d(c_p-3))}{16\pi_d^2} + \frac{4c_p^3A_2}{(1+c_p)^3\pi_u}\right). \tag{A.21}$$

Comparing (A.17) and (A.20), it follows that $\Pi^M > \Pi^L$ if and only if $(64c_p^3\pi_d^2A_2 + (1+c_p)^3(\delta - 2)\pi_uA_3)\pi_u\alpha^2(1+c_p)^3 > 16(1+c_p)^6\pi_d^2\pi_u^2\alpha^2A_4$, which is satisfied if and only if $c_p > 1/3$ and $\delta(1+c_p) > 2(1-c_p)$. This completes the proof. ■

Proof of Proposition 3: By Proposition 2, when $c_p > 1/3$ and $\delta > \frac{2(1-c_p)}{1+c_p}$ are satisfied, then $p^* = p^M$ and $p_s^* = p_s^M$. Substituting (A.15) and (A.16) into (A.11), (A.12), and (A.13) and then subsequently into (13), we obtain

$$\begin{aligned} W^M \triangleq W = {} & \frac{3(1-c_p)^2}{8} + \frac{1}{\alpha}\left(\frac{c_p\delta}{4\pi_d} + \frac{c_p^2(1-c_p)(3-c_p)}{\pi_u(1+c_p)^3}\right) + \\ & \frac{192(3-c_p)c_p^3\pi_d^2A_5 - \pi_u(2-\delta)(1+c_p)^3A_6}{32\pi_u^2\pi_d^2(1+c_p)^7\alpha^2} + O\left(\frac{1}{\alpha^3}\right). \end{aligned} \tag{A.22}$$

where $A_5 = 5c_p^3 - 4c_p^2 + 5c_p - 2$ and $A_6 = 8\pi_dc_p(1-c_p)(7+3c_p^2-14c_p) + \delta\pi_u(1+c_p)^4$.

On the other hand, if it is not the case that both are $c_p > 1/3$ and $\delta > \frac{2(1-c_p)}{1+c_p}$ are satisfied, then, by Proposition 2, $p^* = p^L$ and $p_s^* = p_s^L$. Substituting (A.18) and (A.19) into (A.5), (A.6) and (A.7) and then subsequently into (13), we obtain

$$\begin{aligned} W^L \triangleq W = \frac{3(1-c_p)^2}{8} + \frac{1}{\alpha}\left(\frac{c_p\delta}{4\pi_d} + \frac{c_p^2(1-c_p)(3-c_p)}{\pi_u(1+c_p)^3}\right) + \\ \frac{1}{\alpha^2}\left(\frac{\delta A_7}{32\pi_u\pi_d^2(1+c_p)^4} + \frac{6c_p^3(3-c_p)A_5}{\pi_u^2(1+c_p)^7}\right) + O\left(\frac{1}{\alpha^3}\right) . \end{aligned} \tag{A.23}$$

where $A_7 = \pi_u(1+c_p)^4(\delta-2) - 32c_p^2\pi_d(5-c_p)(1-c_p)$. However, comparing (A.22) and (A.23), it follows that $W^M > W^L$ if and only if $\delta > \frac{2(7-14c_p+3c_p^2)}{(7-c_p)(1+c_p)}$. For $c_p \in (0,1]$, note that $\frac{2(7-14c_p+3c_p^2)}{(7-c_p)(1+c_p)} < \frac{2(1-c_p)}{1+c_p}$ is always satisfied and $\frac{2(7-14c_p+3c_p^2)}{(7-c_p)(1+c_p)} < 1$ is satisfied whenever $c_p > (17-4\sqrt{15})/7$.

By Region III of Lemma 1, (A.15), and (A.16), $p^M \geq p_s^M/\delta - c_p$ is always satisfied and $p_s^M < \delta c_p/(1-\delta)$ is always satisfied for sufficiently high δ. Also, by (A.15) and (A.16), $p^M < p_s^M$ is satisfied if and only if

$$\frac{\delta + c_p(2+\delta) - 2}{4} + \frac{\delta(1+c_p(2\delta-3))}{16c_p\pi_d\alpha} + \frac{3 + c_p(2\delta + c_p - 5 - 2c_p\delta + c_p^2(4\delta-7))}{2(1+c_p)^3\pi_u\alpha} + O\left(\frac{1}{\alpha^2}\right) > 0 . \tag{A.24}$$

Hence, there exists $\underline{c_p}$, $\underline{\eta} > 0$ such that if $\underline{c_p} < c_p < 1/3$ and $\underline{\eta} < \delta < 1$, then, by Proposition 2, $p^* = p^L$ and $p_s^* = p_s^L$, but (A.24) is satisfied; hence, p^M and p_s^M can induce Region III of Lemma 1, which completes the proof. ■

Proof of Proposition 4: By (1) and (2), for sufficiently small α, $c_p > \pi_u u(\sigma)\alpha v$ is satisfied for all $v \in \mathcal{V}$. Hence, $\sigma^*(v) \neq (OP, P)$ for all $v \in \mathcal{V}$, i.e., $v_p = 1$. Because $\delta < 1$ and, by (1), the security risk facing users is $O(\alpha)$, it is simple to establish that the only possible equilibrium consumer market structure characterization under optimal pricing is either $0 < v_d < v_u < 1$ or $0 < v_u' < 1$. Note that when the SaaS price is set to $p_s = \delta v_u'$, the former consumer market structure replicates the latter. Thus, we can focus attention on $0 < v_d < v_u < 1$ and examine the pricing problem in (8).

Using analysis similar to that used in the proof of Lemma 1, the equilibrium equations are given by

$$v_u = v_d + \frac{\delta v_d - p_s}{\delta v_d \pi_d \alpha} , \tag{A.25}$$

and

$$v_d = \sup\left\{v_d \,\middle|\, \delta\pi_d\left(p_s^2 - p_s v_d(1+\delta) + \delta v_d^2(1+\pi_d\alpha(v_d - p))\right) + \right. \\ \left. \pi_u\left(p_s - \delta v_d(1+v_d\pi_d\alpha)\right)\left(p_s + \delta v_d(\pi_d\alpha(1-v_d)-1)\right) = 0\right\}. \tag{A.26}$$

By (A.26), for sufficiently small α, v_d satisfies

$$v_d = \frac{p_s}{\delta} - \frac{\pi_d\alpha p_s(p_s - p\delta)}{\delta^2(1-\delta)} + O(\alpha^2). \tag{A.27}$$

Substituting (A.27) into (A.25) and both expressions subsequently into (7), we obtain

$$\Pi(p,p_s) = \frac{2pp_s\delta + p\delta(1-p-\delta) - p_s^2}{\delta(1-\delta)} - \frac{\pi_d\alpha(p_s - p\delta)^2(2p_s\delta - p\delta - p_s)}{\delta^2(1-\delta)^3} - \\ \frac{\delta^2\pi_u\alpha(p-p_s)^2(1+p_s-p-\delta)}{\delta^2(1-\delta)^3} + O(\alpha^2). \tag{A.28}$$

Differentiating (A.28), the interior-maximizing prices satisfy

$$p^* = \frac{1}{2} - \frac{\pi_u\alpha}{8} + \frac{\pi_u^2\alpha^2}{16(1-\delta)} + O(\alpha^3), \tag{A.29}$$

and

$$p_s^* = \frac{\delta}{2} - \frac{\delta\pi_d\pi_u\alpha^2}{16(1-\delta)} + O(\alpha^3). \tag{A.30}$$

Substituting (A.29) and (A.30) into (A.25) and (A.26) verifies $0 < v_d < v_u < 1$ under optimal pricing. Substituting (A.29) and (A.30) into (A.28) gives the corresponding profits

$$\Pi(p^*,p_s^*) = \frac{1}{4} - \frac{\pi_u\alpha}{8} + \frac{\pi_u^2\alpha^2}{64(1-\delta)} + O(\alpha^3), \tag{A.31}$$

which, by feasibility of $p_s = \delta v_u'$, exceed those obtainable under the $0 < v_u' < 1$ consumer market structure characterization. This completes the proof. ■

Proof of Proposition 5: For the benchmark measures, we refer to August and Tunca (2006) and the benchmark equilibrium characterization therein (see Lemma 1). Thus, for the benchmark case, under high α, v_u is given by the largest root of the polynomial equation

$$\pi_u\alpha v_u^3 + (1 - \pi_u\alpha(p+c_p))v_u^2 - 2pv_u + p^2 = 0. \tag{A.32}$$

By (A.32), for sufficiently large α, v_u satisfies

$$v_u = p + c_p - \frac{c_p^2}{\pi_u \alpha (p+c_p)^2} + \frac{2pc_p^3}{\pi_u^2 \alpha^2 (p+c_p)^5} + O\left(\frac{1}{\alpha^3}\right) . \tag{A.33}$$

Maximizing $p(1-v_u)$, we obtain

$$p_{BM} = \frac{1-c_p}{2} + \frac{2c_p^2(3c_p-1)}{\pi_u\alpha(1+c_p)^3} - \frac{16c_p^3(8c_p^3-5c_p^2+8c_p-3)}{\pi_u^2\alpha^2(1+c_p)^7} + O\left(\frac{1}{\alpha^3}\right) . \tag{A.34}$$

Using (A.34) and similar analysis for the benchmark case, the respective optimal profit and welfare expressions are given by

$$\Pi_{BM} = \frac{(1-c_p)^2}{4} + \frac{2c_p^2(1-c_p)}{\pi_u\alpha(1+c_p)^2} + \frac{4c_p^3(5c_p^3-2c_p^2+5c_p-4)}{\pi_u^2\alpha^2(1+c_p)^6} + O\left(\frac{1}{\alpha^3}\right) \tag{A.35}$$

and

$$W_{BM} = \frac{3(1-c_p)^2}{8} + \frac{c_p^2(3-c_p)(1-c_p)}{\pi_u\alpha(1+c_p)^3} - \frac{6c_p^3(5c_p^4-19c_p^3+17c_p^2-17c_p+6)}{\pi_u^2\alpha^2(1+c_p)^7} + O\left(\frac{1}{\alpha^3}\right) . \tag{A.36}$$

By the proof of Proposition 2, (A.17), (A.20), and (A.35), it follows that

$$\frac{\Pi^* - \Pi_{BM}}{\Pi_{BM}} = \frac{c_p\delta}{\pi_d\alpha(1-c_p)^2} + O\left(\frac{1}{\alpha^2}\right) . \tag{A.37}$$

Similarly, by the proof of Proposition 3, (A.22), (A.23), and (A.36), we obtain

$$\frac{W^* - W_{BM}}{W_{BM}} = \frac{2c_p\delta}{3\pi_d\alpha(1-c_p)^2} + O\left(\frac{1}{\alpha^2}\right) , \tag{A.38}$$

which completes the proof. ■

Proof of Proposition 6: Similar to the proof of Proposition 5, for the benchmark case, under low α, v_u is given by

$$v_u = -\frac{1-\pi_u\alpha}{2\pi_u\alpha} + \frac{1}{2\pi_u\alpha}\sqrt{(1-\pi_u\alpha)^2 + 4\pi_u\alpha p} \tag{A.39}$$

By (A.39), for sufficiently low α, v_u satisfies

$$v_u = p + \pi_u p(1-p)\alpha + \pi_u^2 p(1-p)(1-2p)\alpha^2 + O\left(\alpha^3\right) . \tag{A.40}$$

Maximizing $p(1-v_u)$, we obtain

$$p_{BM} = \frac{1}{2} - \frac{\pi_u \alpha}{8} + \frac{\pi_u^2 \alpha^2}{16} + O\left(\alpha^3\right) . \tag{A.41}$$

Using (A.41), the respective optimal profit and welfare expressions are given by

$$\Pi_{BM} = \frac{1}{4} - \frac{\pi_u \alpha}{8} + \frac{\pi_u^2 \alpha^2}{64} + O\left(\alpha^3\right) . \tag{A.42}$$

and

$$W_{BM} = \frac{3}{8} - \frac{\pi_u \alpha}{4} + \frac{5\pi_u^2 \alpha^2}{128} + O\left(\alpha^3\right) . \tag{A.43}$$

By the proof of Proposition 4, (A.31), and (A.42), it follows that

$$\frac{\Pi^* - \Pi_{BM}}{\Pi_{BM}} = \frac{\delta \pi_u^2 \alpha^2}{16(1-\delta)} + O\left(\alpha^3\right) . \tag{A.44}$$

By (10), (A.25), (A.29), and (A.30), we obtain

$$\frac{W^* - W_{BM}}{W_{BM}} = \frac{5\delta \pi_u^2 \alpha^2}{48(1-\delta)} + O\left(\alpha^3\right) . \tag{A.45}$$

By (A.44) and (A.45), the result follows. ■

Proof of Proposition 7: By Proposition 2, if $c_p > 1/3$ and $\delta > \frac{2(1-c_p)}{1+c_p}$ are satisfied, then by substituting (A.11), (A.12), and (A.13) into (12) and subsequently into (21), we obtain

$$\hat{SL}^* = c_p + \frac{\delta - c_p(4-\delta)}{4\pi_d(1-c_p)\alpha} + O\left(\frac{1}{\alpha^2}\right) . \tag{A.46}$$

If either $c_p > 1/3$ or $\delta > \frac{2(1-c_p)}{1+c_p}$ is not satisfied, then making analogous substitutions into (15) and (21), it follows that in this case $\hat{SL}^*$ also satisfies (A.46). On the other hand, using a similar train of logic for the benchmark case, it can be shown that

$$\hat{SL}_{BM} = c_p + O\left(\frac{1}{\alpha^2}\right) . \tag{A.47}$$

By (A.46) and (A.47), it follows that

$$\frac{\hat{SL}^{*} - \hat{SL}_{BM}}{\hat{SL}_{BM}} = \frac{\delta - c_p(4-\delta)}{4c_p\pi_d(1-c_p)\alpha} + O\left(\frac{1}{\alpha^2}\right) , \tag{A.48}$$

which proves the result. ■

Proof of Proposition 8: Because $c_p < 1/3$, part (ii) of Proposition 2 applies, hence by (14) and (16), consumer surplus is given by $CS = W - \Pi$. Following the proofs of Proposition 2 and Proposition 3, by (A.20) and (A.23), we obtain

$$\begin{aligned} CS^{*} = & \frac{(1-c_p)^2}{8} + \frac{c_p^2(1-c_p)(1-3c_p)}{\pi_u\alpha(1+c_p)^3} + \frac{\delta}{32\pi_d^2\alpha^2}\left(2 - \delta - \frac{64c_p^2\pi_d(1-c_p(4-c_p))}{\pi_u(1+c_p)^4}\right) \\ & - \frac{2c_p^3(25c_p^4 - 51c_p^3 + 57c_p^2 - 49c_p + 10)}{\pi_u^2\alpha^2(1+c_p)^7} + O\left(\frac{1}{\alpha^3}\right) . \end{aligned} \tag{A.49}$$

Similarly, by the proof of Proposition 5, (A.35), and (A.36), we obtain

$$CS_{BM} = \frac{(1-c_p)^2}{8} + \frac{c_p^2(1-c_p)(1-3c_p)}{\pi_u\alpha(1+c_p)^3} - \frac{2c_p^3(25c_p^4 - 51c_p^3 + 57c_p^2 - 49c_p + 10)}{\pi_u^2\alpha^2(1+c_p)^7} + O\left(\frac{1}{\alpha^3}\right) . \tag{A.50}$$

Comparing (A.49) and (A.50), it follows that

$$CS^{*} - CS_{BM} = \frac{\delta}{32\pi_d^2\alpha^2}\left(2 - \delta - \frac{64c_p^2\pi_d(1-c_p(4-c_p))}{\pi_u(1+c_p)^4}\right) + O\left(\frac{1}{\alpha^3}\right) , \tag{A.51}$$

which proves the result. ■

Proof of Proposition 9: Technically, we will prove that (i) there exist $\overline{\gamma} \in (0,1)$ and $\overline{\omega} > 0$ such that if $\gamma < \overline{\gamma}$, $\alpha < \overline{\omega}$, and $\pi_u' = \gamma\pi_u$, then $\hat{SL}_{BM}(\cdot \mid \pi_u') < L^{*}(\cdot \mid \pi_u)$ and $W_{BM}(\cdot \mid \pi_u') > W^{*}(\cdot \mid \pi_u)$, and (ii) there exist $\overline{\beta} \in (0,1)$ and $\underline{\omega} > 0$ such that if $\beta < \overline{\beta}$, $\alpha > \underline{\omega}$, and $c_p' = \beta c_p$, then $\hat{SL}_{BM}(\cdot \mid c_p') < L^{*}(\cdot \mid c_p)$ and $W_{BM}(\cdot \mid c_p') > W^{*}(\cdot \mid c_p)$.

For part (ii), by the proofs of Proposition 3 and 5, (A.22), (A.23), and (A.36), it follows that

$$W_{BM}(\cdot \mid c_p') > W^{*}(\cdot \mid c_p) \iff (1-\beta c_p)^2 < (1-c_p)^2 + O\left(\frac{1}{\alpha}\right) , \tag{A.52}$$

which is satisfied for sufficiently small β. For part (i), by the proof of Proposition 4, (10), (A.25), (A.29), and (A.30), we obtain

$$W^* = \frac{3}{8} - \frac{\pi_u \alpha}{4} + \frac{5\pi_u^2 \alpha^2}{128(1-\delta)} + O\left(\alpha^3\right) \ . \tag{A.53}$$

Similarly, by the proof of Proposition 6 and (A.43), it follows that $W_{BM}(\cdot \,|\, \pi'_u) > W^*(\cdot \,|\, \pi_u)$ is satisfied whenever

$$\begin{aligned} & W_{BM}(\cdot \,|\, \pi'_u) > W^*(\cdot \,|\, \pi_u) \\ \iff \quad & \tfrac{3}{8} - \tfrac{\gamma \pi_u \alpha}{4} + O\left(\alpha^2\right) > \tfrac{3}{8} - \tfrac{\pi_u \alpha}{4} + O\left(\alpha^2\right) \\ \iff \quad & \gamma < 1 - O\left(\alpha\right) \ , \end{aligned} \tag{A.54}$$

which completes the proof for the welfare comparisons. The analysis for the per-user security losses follows a similar logic and has been omitted. ■